Planning for the Unplanned: Oracle® 10*g*™ R2 Disaster Recovery with a NetApp Storage System

SnapMirror® Async and Sync

Jawahar Lal, NetApp

TR-3455

Table of Contents

1. DISASTER RECOVERY: EXECUTIVE SUMMARY

As the need for uninterrupted availability of enterprise data is growing, more businesses are striving for 24x7 system availability and can't afford any downtime. In this era of continuous information availability, the complete and rapid recovery from a disaster is not nice to have, but is a necessity.

Whether you're a business, a government agency, a healthcare organization, or an educational institution, you must ensure that you're prepared for when, not if, disaster strikes. A protracted interruption in your organization's ability to access data will disrupt business operations. This can lead to the loss of customers and revenue, a drop in share price, or possible noncompliance fines for failing to protect and/or provide information promptly when required.

A well-planned, rehearsed, and tested Disaster Recovery solution can save time and money for your organization by offering small recovery windows with no or acceptable data loss. NetApp offers an array of proven, low-cost, and simplified data protection and Disaster Recovery solutions for your organization's data.

2. INTRODUCTION

As recent world events have proven, disaster can strike anywhere and anytime. The question you need to answer is how well prepared your organization is to recover in the event of a disaster. Though most business houses, even the small ones, are meticulous when it comes to thinking ahead and devising future strategies to take the company forward, a common mistake is ignoring the possibility of a disaster crippling the organization. Critical data loss could well be the fatal prescription for an organization that is otherwise doing well.

Survey after survey shows they're in the majority: Over 50% of companies make no effort whatsoever to prevent avoidable disasters or to put into place strategies for recovering from outage events that can't be avoided. Of those companies that do plan, fewer than 50% actually test the strategy they develop, which is like having no strategy at all.

From time to time different groups, including professional accounting organizations, universities, and the U.S. government, study the results of disasters and downtime on business and employment. The results are never good. Among the most-published results are:

93% of companies that suffer a significant data loss are out of business within five years (U.S. Bureau of Labor).
43% of U.S. businesses never reopen after a disaster, and 29% close within two years (University of Wisconsin).
30% of computer users say they spend the equivalent of one week per year reconstructing lost data (3M Corporation).

All these statistics suggest that organizations must plan for Disaster Recovery up front. The Disaster Recovery Plan is a corporate survival kit in the eventuality of disaster. The continued operations of an organization depend on management's awareness of potential disasters, the ability to develop a Disaster Recovery Plan to minimize disruptions of critical functions, and the capability to recover and restore vital business operations successfully and quickly.

A Disaster Recovery Plan is a comprehensive statement of consistent actions to be taken before, during, and after a disaster. The plan should be documented and tested to ensure the continuity of operations and availability of critical resources in the event of a disaster. The primary objectives of the plan are to protect the organization in the event that all or part of its IT services are rendered unusable, minimize the disruption to critical business operations, ensure some level of organizational stability, and provide an orderly recovery after a disaster.

Incidents like 9/11 and Hurricane Katrina give businesses a chance to see their Disaster Recovery Plan in action. While some companies pass with flying colors, the plans of others are exposed as incomplete, unrealistic, and technologically flawed. Those companies with untested or poorly tested plans will eventually discover that they aren't as protected as they thought they were.

There are several approaches to protect data and maintain data availability in the face of hardware, software, or even site failures. Backups provide a way to recover lost data from an archival medium (tape or disk). Redundant hardware technologies also mitigate the damage caused by hardware failures. Data replication or mirroring provides a third mechanism to ensure data availability and minimize downtime. The Network

Appliance™ SnapMirror product provides a fast and flexible enterprise solution for data replication over local area, wide area, and Fibre Channel (FC) networks. If a disaster strikes at a source site, the replicated mission-critical data can immediately be made available at a remote site, ensuring uninterrupted operation and data availability. For detailed information on popular Disaster Recovery solutions offered by NetApp, please visit the data protection solutions page on the NetApp Web site (http://www.netapp.com/solutions/data_protection.html).

3. PURPOSE AND SCOPE

The scope of this document is limited to disaster recovery of an Oracle 10*g* R2 single-instance database. The two disaster recovery solutions covered in this document are:

1. Disaster Recovery using SnapMirror Async
2. Disaster Recovery using a mix of SnapMirror Async and Sync

Using SnapMirror, it is now possible to recover from a disaster at a remote physical location. If critical data is mirrored to a different physical location, a serious disaster no longer means prolonged data unavailability. The mirrored data can be made available to clients across the network until the damage caused by the disaster is repaired. Recovery may include recovery from corruption, natural disaster at the source site, accidental deletion, sabotage, etc. SnapMirror also allows application server layer information to be replicated to the Disaster Recovery site. In the event of disaster, once the Disaster Recovery site is operational, all applications can be switched over to the servers at the Disaster Recovery site and all application traffic can be directed to these servers for as long as necessary to recover the source site. Once the source site is recovered, SnapMirror can be used to transfer the data efficiently back to the primary site. After the production site takes over normal application operation again, SnapMirror transfers to the Disaster Recovery site can resume without requiring a second complete data transfer.

4. UNDERSTANDING DISASTER RECOVERY LANGUAGE

1) Business Continuity Plan (BCP)

Business Continuity Plan (BCP) describes processes and procedures an organization puts in place to ensure that essential business processes continue during and after a disaster. Normally it takes into account the protection of the whole organization, including buildings, IT infrastructure, employees, and all other resources. The main objective of BCP is to prevent interruption of mission-critical services and to reestablish full functions as swiftly and smoothly as possible.

2) Disaster Recovery Plan (DRP)

Disaster Recovery Plan is a subset of the Business Continuity Plan and focuses solely on the protection and recovery of the mission-critical data and the IT infrastructure. It details step-by-step procedures and processes for the IT staff to follow during and after the disaster. It describes the recovery priority and objectives (RPO and RTO) for each application.

3) Recovery Time Objective (RTO)

The Recovery Time Objective (RTO) indicates the time spent in bringing the application up and resuming the operation after the disaster. It is also known as acceptable downtime after the disaster. The unit of measure for RTO is time, with values ranging from seconds to days or weeks. Lower the application's RTO value greater the organization's dependence on that particular application, and consequently higher the recovery priority after the disaster.

4) Recovery Point Objective (RPO)

The Recovery Point Objective (RPO) defines data currency. The unit of measure for the RPO is also time, with values ranging from seconds to days or weeks. It denotes how current the data should be after the recovery. In another words, it defines acceptable data loss from the point the disaster starts. Lower an application's RPO value, greater the organization's dependence on that particular process, and consequently higher the priority when recovering the systems after the disaster.

5) Disaster Tolerance

Greater awareness of the need for Disaster Recovery is prompting application architects to build disaster readiness into the business systems they design. Disaster tolerance is a term used to signify a system with some ability to withstand major disruption. Several technologies are used to provide disaster tolerance, including hardware redundancy, data replication, server clustering, and remote data centers.

6) High Availability (HA)

High availability is an architecture that maximizes the data availability. It is a subcategory of Disaster Recovery. The ultimate disaster-tolerant system is classed as a high-availability (HA) system. The HA systems are designed to eliminate application downtime by using redundant hardware and networking components and specialized application and operating system software. HA systems can seamlessly route around failures in the computing infrastructure without affecting end-user access to data.

The resilience of this system is often measured in terminology borrowed from the telecommunications industry. For example, a configuration that offers 99.999% availability (also known as five nines), can have only 5 minutes of planned or unplanned downtime in any a given year.

7) Archive Logging

Archive logging is a database feature that enables retention of the transaction logs. The retained transaction logs are called archive logs. Using archive and active logs, a database roll-forward recovery is possible to any point in time before the failure occurred, rather than only to the point in time of a full backup. The archived logs can be moved off line and still be used for roll-forward recovery.

8) Media Recovery

Media recovery is a user-initiated data recovery. It can be used to recover out-of-date or damaged datafiles, SPFile, or control files. The recovery is performed by applying archive logs followed by active logs. The various scenarios that require media recovery are:

Restore datafiles from backup.
Restore control files from backup.
Datafiles are taken offline without OFFLINE NORMAL option.
Datafiles are out of date with the corresponding control file.

A database can't be opened if datafiles need media recovery.

9) Instance or Crash Recovery

Instance crash recovery process is a special form of recovery, which happens when a database instance is started for the first time after a crash (or SHUTDOWN ABORT). The crash recovery uses only online transaction logs and the goal is to bring the datafiles to a transaction-consistent state, preserving all committed changes up to the point when the instance failed.

10) Application-Coordinated Snapshot™ Copy

An application-coordinated Snapshot copy is a Snapshot copy created manually after putting a database in hot backup mode, it guarantees the database consistency. That means a database recovery is guaranteed from an application-coordinated Snapshot copy.

5. SNAPMIRROR: A QUICK OVERVIEW

In order to protect your organization's data from disaster and ensure quick and smooth recovery, your data needs to be replicated to one or more other physical locations. NetApp SnapMirror technology allows data replication between two NetApp storage systems. The NetApp storage system from which data is transferred is referred to as the SnapMirror source, and the NetApp storage system to which the data is transferred is referred to as the SnapMirror destination. The SnapMirror source and destination can be miles apart, provided that both NetApp storage systems can communicate with each other across a network.

NetApp SnapMirror technology supports Volume SnapMirror (VSM) as well as qtree SnapMirror (QSM). SnapMirror source volumes and qtrees are writable data objects, but the SnapMirror destination volumes and qtrees are read-only, usually on a separate storage system. In the case of a disaster where the source volumes or qtrees go down, the replicated data at the destination volumes or qtrees can be made available by making the volumes or qtrees writable. The SnapMirror configuration details are maintained in a configuration

file called `snapmirror.conf`, which resides on the destination NetApp storage system. This file, along with information entered via the `snapmirror.access` option or the `snapmirror.allow` file, is used to establish relationships between specified source volumes or qtrees and the destination volume or qtree where the mirrored data is kept.

NetApp Data ONTAP® supports SnapMirror Async as well as SnapMirror Sync. In the SnapMirror Async mode, SnapMirror performs incremental, block-based replication as per the frequency defined in `snapmirror.conf` file. Write requests to the SnapMirror source are acknowledged as soon as they are written to its NVRAM and are not delayed until the SnapMirror destination has received and/or processed the request. Performance impact on the SnapMirror source filer is minimal, as long as the system is configured with sufficient CPU and disk I/O resources. The first and most important step in Async mode involves the creation of a one-time baseline transfer of the entire data set from the SnapMirror source system to the SnapMirror destination system. This is required before incremental updates can be performed. After the baseline transfer is complete, scheduled or manually triggered updates can occur. Each update transfers only the new and changed blocks from the source to the destination storage system. Because asynchronous replication is periodic, SnapMirror is able to consolidate writes and conserve network bandwidth, thereby minimizing the impact on write throughput and write latency.

NetApp Data ONTAP supports SnapMirror Sync to meet the very high availability requirements for certain environments where all data changes written to a production site must be replicated to a remote site synchronously. SnapMirror in Synchronous or Sync mode is a mode of replication that sends updates from the source volumes or qtree to the destination volumes or qtree as they occur, rather than according to a predetermined schedule. This guarantees that data written on the source system is protected on the destination even if the entire source system goes down. With synchronous mode, each time a transaction attempts to write data to disk, the data is sent to both the source and destination storage systems in parallel. It is not until both NetApp storage systems have committed the data associated with the write operation to NVRAM that the system acknowledges that the transaction is complete. In other words, the application that initiated the write operation must wait until it receives the acknowledgement from both the source and destination storage systems before it can continue.

SnapMirror Semi-Sync mode is a variation of SnapMirror Sync mode and it provides a middle ground that keeps the source and destination file systems more closely synchronized than the asynchronous mode, but with less impact on application performance. Configuration of Semi-Synch mode is nearly identical to the configuration of Sync mode, with the exception being the addition of an option that specifies how many writes, seconds, or ops can be outstanding (unacknowledged by the destination system) before the source system delays acknowledging write operations from clients. Internally, Semi-Sync mode works identically to Sync mode in most cases. The only difference lies in how quickly client writes are acknowledged; the replication methods used are the same.

Note: The Sync and Semi-Sync modes are supported only with VSM.

For more details on SnapMirror deployment and Implementation, please refer to Technical Reports "SnapMirror Deployment and Implementation Guide" and "Synchronous SnapMirror Design and Implementation Guide" on NOW™ (NetApp on the Web) at http://now.netapp.com.

6. REQUIREMENTS AND ASSUMPTIONS

1) General Assumptions

In order to take maximum advantage of the procedures described in this document, it is assumed that readers of this document are familiar with the following:

- Commands and operations of Data ONTAP and NetApp storage system
- Administration and operation of Oracle 10*g* R2 instance and database
- UNIX® system administration commands
- NetApp SnapMirror technology
- Disaster Recovery concept

The NetApp storage systems used to produce this document are loaded with Data ONTAP 7G and are licensed for NFS, FCP, iSCSI, SnapMirror, and SnapMirror Sync products.

It is also assumed that the UNIX hosts used for the production database have Oracle 10*g* R2 software installed and a single database instance created.

In order to produce this document we used NFS protocol. If you are using a SAN environment, then make sure that database hosts have the following products installed and configured:

Appropriate HBA
SanSurfer Utility (installed if HBA used is from Qlogic)
NetApp Host Attach/Support Kit

Please check the Compatibility and Configuration Guide for NetApp FCP and iSCSI Products to find out about supported HBAs.

2) Environment Assumptions

This document covers Disaster Recovery solutions offered by NetApp for Oracle 10*g* R2 single database instance using SnapMirror Async as well as Sync technology. The scripts and process steps contained in this document may require significant modifications to run under your version of UNIX. The sample scripts in this document assume the following:

The NetApp storage system used for the SnapMirror source is `'ntapsrc'`
The NetApp storage system used as the SnapMirror destination is `'ntapdst'`
The aggregate on the NetApp storage systems used for database storage is `'dbaggr'`
The flexible volume used to store database data is `'dbdata'`
The flexible volume used to store database transaction logs is `'dblogs'`
The flexible volumes `dbdata` and `dblogs` reside in aggregate `dbaggr`
At the database host, the mount points used are `/mnt/dbdata` and `/mnt/dblogs`
Oracle Home resides on a Linux® (RHEL 4) database host

It is also assumed that the database host used for accessing the database at the Disaster Recovery site has a similar setup to the database host on the primary site and has all required privileges to access the volumes on the SnapMirror destination NetApp storage system.

3) Security and Access Issues

You need to make sure that the FlexVol™ volumes to be used for the database's data and transaction logs have their security style set appropriately. If the database host used is a UNIX host, then the security style must be set to 'UNIX'. The security style can be changed by executing the following command on the NetApp storage system:

```
qtree security <volume name> unix
```

For example, to change the security style for a volume named `dbdata`, you would execute the following command on the NetApp storage system:

```
qtree security /vol/dbdata unix
```

4) Network Connectivity

You also need to make sure that the database hosts and NetApp storage systems can communicate with each other through the network. This is done by making appropriate entries to the `/etc/hosts` files on the NetApp storage system as well as on the database host systems.

Add the following line to the database host's `/etc/hosts` file if it doesn't already exist:

```
<NetApp storage system IP>  <NetApp storage system name>
```

For example, to add an entry to the `/etc/hosts` file on the database server for the NetApp storage system named `ntapsrc` that has IP address `10.32.70.134`, you would add the following line:

```
10.32.70.134 ntapsrc
```

Add the following line to the `/etc/hosts` file on a NetApp storage system if it doesn't already exist:

```
<database server IP>  <database server name>
```

For example, to add an entry to the `/etc/hosts` file on the NetApp storage system for a database server named `dbhost1` that has IP address `172.32.70.43`, you would add the following line:

```
172.32.70.43 dbhost1
```

5) Mount the NetApp storage system's root volume to the database server

In order to mount the NetApp storage system's `root` volume to the database host and make necessary changes to some configuration files, the user `root` on the database sever must have access to the root volume `/vol/vol0`. For example, to grant access on a root volume named `/vol/vol0` on a NetApp storage system to the user `root` on the database host system (`dbhost1`), you would execute the following command on the storage systems:

```
exportfs –p rw=dbhost1,root=dbhost1,anon=0 /vol/vol0
exportfs –a
```

6) Enable `rsh` access from the database server

If you intend to use `'rsh'` commands from your database host, then add the IP address of the database host to the `/etc/hosts.equiv` file on the NetApp storage system. The entry should look similar to the following:

```
<hostsrc IP>
```

7) Required permissions on the volumes to be used for the database

Before you create a database on the FlexVol volumes on the NetApp storage system, you need to mount them to a database host system. In order to mount FlexVol volumes, the user `root` on the database host must have access to them. Execute the following commands on the NetApp storage systems to grant access on FlexVol volumes:

```
exportfs -p rw=<db host name>,root=<db host name>,anon=0 <volume name>
```

For example, to grant access on the FlexVol named `dbdata` to the user `root` on the database server named `dbhost1`, you would execute the following command:

```
exportfs -p rw=172.17.38.112,root=172.17.38.112,anon=0 /vol/dbdata
exportfs –a
```

Grant permission on all FlexVol volumes to be used for the database using the above command.

8) Mount and change ownership of the file system on the mount point

The FlexVol volumes to be used for the database need to be mounted on the database host system by executing the following command:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600 <Netapp
storage system name>:<volume name> <mount point>
```

For example, to mount a volume named `dbdata` on a mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
<ntapsrc>:/vol/dbdata /mnt/dbdata
```

To install Oracle and create a database successfully, you need to make sure that the file system on the mount points is owned by the user `'oracle'` on the database server. Change the ownership of the mounted volumes to the user `oracle` by executing the following command at the database server:

```
chown –R oracle:dba <mount point>
```

For example, to change the ownership of a file system on the mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
chown –R oracle:dba /mnt/dbdata
```

7. NETWORK AND STORAGE INFRASTRUCTURE

The following network diagram shows a very simple and basic architecture for the Oracle database Disaster Recovery scenario using a NetApp storage system at the back end and that was used to produces this document.

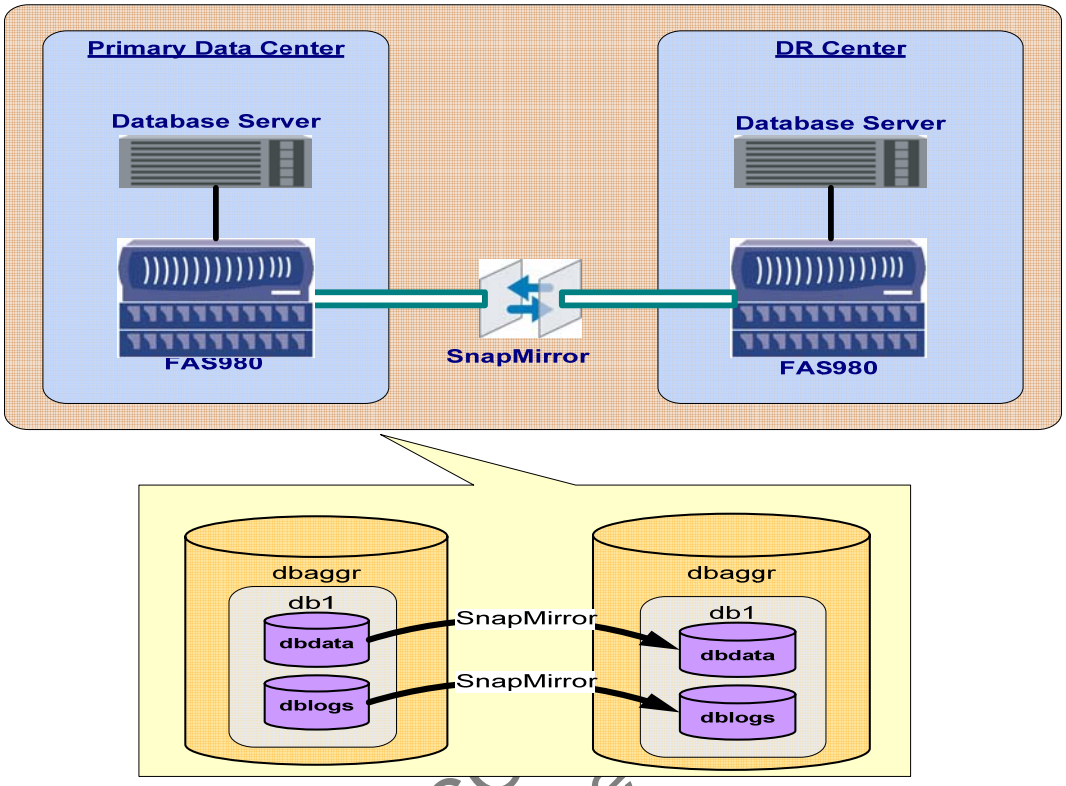

Figure 1) Network and storage infrastructure.

In this architecture, the Oracle Home resides on the database server. The datafiles, redo logfiles, and archive logfiles for the database reside on the FlexVol volumes on a NetApp storage system. The volumes used for the database on the NetApp storage system at the primary data center are replicated to a secondary location using SnapMirror. The secondary location is referred to as the Disaster Recovery site through out this document.

8. CONFIGURATION DETAILS

In order to produce this document we used Oracle Database 10*g* R2, Enterprise Edition. The database host was running 32-bit Red Hat Enterprise Linux Kernel 2.6 on a Sun Fire V20Z dual process machine. For database storage, we used NetApp FAS900 series (FAS980). The environment details are:

- Database – Oracle 10*g* R2, single database instance (orcl)
- Database archive log - Enabled
- Oracle Home – Local on database Host (`/home/oracle/ora10g/product/10.2.0/db_1`)
- NetApp storage system volumes used for the database data and transaction logs are `dbdata` and `dblogs` respectively
- Soft links created to follow the OFA directory structure:
  ```
  ln -s /mnt/dbdata /home/oracle/ora10g/dbdata
  ln -s /mnt/dblogs /home/oracle/ora10g/dblogs
  ```

Oracle Database layout:

Control files:

```
/home/oracle/ora10g/dbdata/orcl/control01.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
/home/oracle/ora10g/dblogs/orcl/control03.ctl
```

Redo logs:

```
/home/oracle/ora10g/dblogs/orcl/redo_01.log
```

```
/home/oracle/ora10g/dblogs/orcl/redo_02.log
/home/oracle/ora10g/dblogs/orcl/redo_03.log
```

Datafiles:

```
/home/oracle/ora10g/dbdata/orcl/*.dbf
```

It is recommended that a control file copy should be stored on every disk drive that stores members of online redo log groups, if the online redo log is multiplexed. By storing control files in these locations, you minimize the risk that all control files and all groups of the online redo log will be lost in a single disk failure. For example, we have placed one control file named `home/oracle/ora10g/`**`dblogs`**`/orcl/control03.ctl` on the same volume where database redo logfiles reside.

9. DATABASE SETUP

As described in section 6, it is assumed that you already have Oracle 10*g* R2 software installed and a single database instance created, and that the database has its data and transaction logs on NetApp storage system volumes. In order to complete the scenario discussed in this document, you need to complete the following steps:

a). Log in to the database and create a table by executing the following command:

```
create table scott.tab1(col1 number not null, col2 character(240));
```

b). Create a database stored procedure by executing the script `demobld.sql`. The script can be found in Appendix A.

```
sql>@demobld.sql
```

The procedure `demobld` inserts rows in the table `scott.tab1` and generates the database load.

10. CONFIGURE SNAPMIRROR

To protect your organization's data using NetApp SnapMirror technology, you need to complete the following very simple SnapMirror configuration process steps:

a). Identify the SnapMirror source and the destination NetApp storage systems. The source NetApp storage system should already have an Oracle 10*g* R2 database created on it. For example, we used a NetApp storage system named `ntapsrc` that has Oracle 10*g* R2 database create on it, as the SnapMirror source and another NetApp storage system named `ntapdst` that is on the Disaster Recovery site as the SnapMirror destination.

b). You need to identify the volumes on the source NetApp storage system that need to be replicated to one or more other NetApp storage systems using SnapMirror.

c). You need to define a schedule for each SnapMirror relationship in the `/etc/snapmirror.conf` file, which resides on the SnapMirror destination's NetApp storage system. The `/etc/snapmirror.conf` file controls where data is replicated from and how often the mirror is updated. The entry in the `/etc/snapmirror.conf` file should be in the following format:

```
<source NetApp storage system>:<vol name | qtree path>  <destination NetApp
storage system>:<vol name | qtree path>  '-' schedule
```

Where:
The `schedule` is defined by `[minute] [hour] [days of month] [days of week]`

A snippet from the `/etc/snapmirror.conf` file:

```
ntapsrc:dbdata ntapdst:dbdata – 0,30 9-17 * 1-5
ntapsrc:dblogs ntapdst:dblogs – 0,30 * * 1-5
```

In the above snippet, the volume named `dbdata` on the NetApp storage system named `ntapsrc` is replicated to the volume named `dbdata` on a NetApp storage system `ntapdst`. The argument dash '-' represents that data is replicated at the fastest rate possible. The NetApp storage system `ntapdst`

updates the volume `dbdata` at every 30-minute interval between 9 a.m. and 5 p.m., Monday through Friday. The asterisk (*) in the example means that the mirror is updated on every day of the month.

d). The SnapMirror feature needs to be enabled SnapMirror on the source as well as on the destination NetApp storage system by executing the following command:

```
options snapmirror on
```

e). In order to complete a one-time baseline transfer, each SnapMirror relationship needs to be initialized by executing the following command from the SnapMirror destination NetApp storage system:

`snapmirror initialize –S <source NetApp storage system>:volume <destination NetApp storage system>:volume`

For example, to initialize the baseline database transfer for the volume named `dbdata`, you would execute the following command on the SnapMirror destination NetApp storage system:

```
snapmirror initialize –S ntapsrc:dbdata ntapdst:dbdata
```

After the baseline transfer, the SnapMirror updates can be triggered based on a predefined schedule in the `/etc/snapmirror.conf` file. It is also possible to trigger a SnapMirror update manually by executing the following command from the destination storage system:

`snapmirror update <volume name>`

For example, you would execute the following command on the SnapMirror destination NetApp storage system to start the manual SnapMirror update for a volume named `dbdata`:

```
snapmirror update dbdata
```


11. DATABASE RECOVERY USING SNAPMIRROR ASYNC

SnapMirror Async transfers the changed data blocks to the SnapMirror destination based on Snapshot copy comparison. On each SnapMirror update, a new Snapshot copy is created at the SnapMirror source NetApp storage system and it is compared against the Snapshot copy from the previous SnapMirror update to determine the changes. After changes are determined, the changed data blocks are transferred to the destination NetApp storage system and the old Snapshot copy gets deleted. The SnapMirror update and application-coordinated Snapshot schedule is summarized in the following table:

Flexible Volume	SnapMirror	App. Coordinated Snapshot
dbdata	**2 hrs**	**30 min**
dblogs	**30 min**	**10 min**

The snippet from the `/etc/snapmirror.conf` file looks somewhat similar to the following:

```
ntapsrc:dbdata ntapdst:dbdata – 0 0-22/2 * 1-5
ntapsrc:dblogs ntapdst:dblogs – 0,30 * * 1-5
```

In order to simulate disaster on the primary site and perform recovery on the Disaster Recovery site, you would need to complete the following steps:

Step 1- Complete the SnapMirror base data transfer and run the load

Initialize the SnapMirror relationship for each volume used for the database and complete one-time baseline transfer by executing the following command on the destination NetApp storage system:

`snapmirror initialize –S <source NetApp storage system>:<volume name> <destination NetApp storage system>:<volume name>`

For example, to initialize the baseline transfer for a volume named `dbdata`, you would execute the following command on the destination NetApp storage system:

```
snapmirror initialize –S ntapsrc:dbdata ntapdst:dbdata
```

Monitor the SnapMirror status by executing the following command on the destination storage system:

```
snapmirror status
```

Let us assume our base line transfer started at 12:00 p.m. and completed at 12:25 p.m.

After creating a table and stored procedure as described in section 9, you need to connect to the database and start running the load by executing the following commands on the database server:

```
connect scott/tiger
execute scott.demobld(25000000);
```

Copy the `do_snap`, `begin_bkup`, and `end_bkup` scripts to a work directory on the database server and start creating application-coordinated Snapshots copies by executing the `do_snap` script on the database server:

```
./do_snap
```

The above command can be executed as a `cron` job from the database server.

Step 2- Simulate the disaster

Let us assume a that the disaster strikes at 3:05 p.m. In order to simulate the disaster, we made the SnapMirror source NetApp storage system unavailable by executing the following command:

```
halt
```

At this point, the system state on the primary site will look somewhat similar to Figure 2.

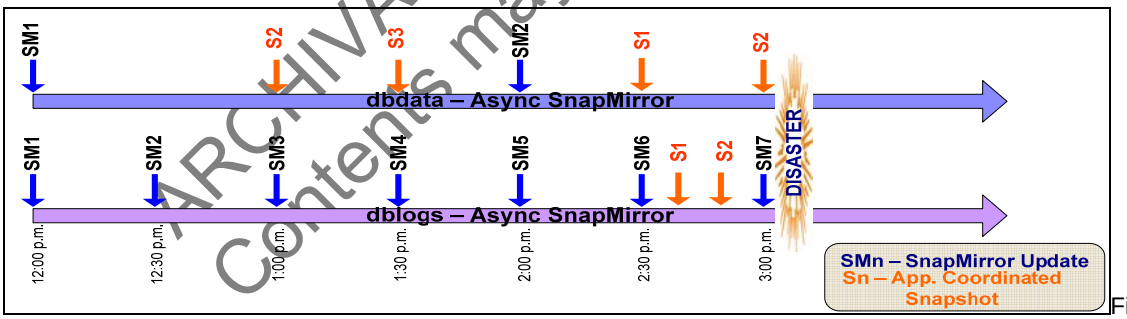

Figure 2) Database state at the primary site at the time disaster struck.

After a disaster, the Disaster Recovery site will be in SnapMirror Snapshot consistent state. That means the volumes at the Disaster Recovery site will have the same data, including application-coordinated Snapshot copies, as the corresponding source volume had at the time of successful SnapMirror update. For our test scenario, `dbdata` and `dblogs` are the SnapMirror destination volumes corresponding to the SnapMirror source volumes used for the database; therefore, these two volumes will be in the following states:

`dbdata` – will be in consistency with the source volume at the time of the SnapMirror update at 2.00 p.m.
`dblogs` – will be in consistency with the source volume at the time of the SnapMirror update at 3.00 p.m.

Step 3 - Recovery

After the initial transfer, the destination volumes are available to clients, but in a read-only state. The status of a destination will show that it is snapmirrored. A disaster makes the source unavailable; to use the destination volumes for writing as well as reading, you would need to end the SnapMirror relationship for each volume used for the database by executing the following command on the destination NetApp storage system:

```
snapmirror break <volume name>
```

For example, to make a volume named `dbdata` writable, you would execute the following command on the NetApp storage system:

```
snapmirror break dbdata
```

This command changes the destination volume's status from **snapmirrored** to **broken-off**, thus making it writable.

After a disaster, the destination volumes will be in SnapMirror Snapshot consistency state with their corresponding source volumes used for the database's data and will have the same data, including application-coordinated Snapshot copies, as the corresponding source had at the time of successful SnapMirror update. The database transaction logs volume is sync snapmirrored, and therefore is in NVLOG or CP consistency state based on the configuration. In this type of configuration, there are four ways of database recovery after the disaster:

1) Recovery from the SnapMirror update points of the `dbdata` and `dblogs` volumes.
2) Recovery from an application-coordinated Snapshot copy of the `dbdata` volume and SnapMirror update point of the `dblogs` volume.
3) Recovery from the SnapMirror updated point of the `dbdata` volume and an application-coordinated Snapshot copy of the `dblogs` volume.
4) Recovery from the application-coordinated Snapshot copies of the `dbdata` and `dblogs` volumes.

1) Recovery from the SnapMirror update point

In this scenario of database recovery, the RPO (Recovery Point Objective) is determined by the time from the last successful SnapMirror update of the transaction logs volume. The data after the last SnapMirror update of the transaction logs volume will be lost. The RTO (Recovery Time Objective) will vary based on the transaction logs that need to be applied to recover the database. For example, in our Disaster Recovery scenario, the disaster struck at 3:05 p.m. and the last SnapMirror update for the `dblogs` volume occurred at 3:00 p.m. therefore, the RPO will be 5 minutes. Only 5 minutes worth of data needs to be reconstructed.

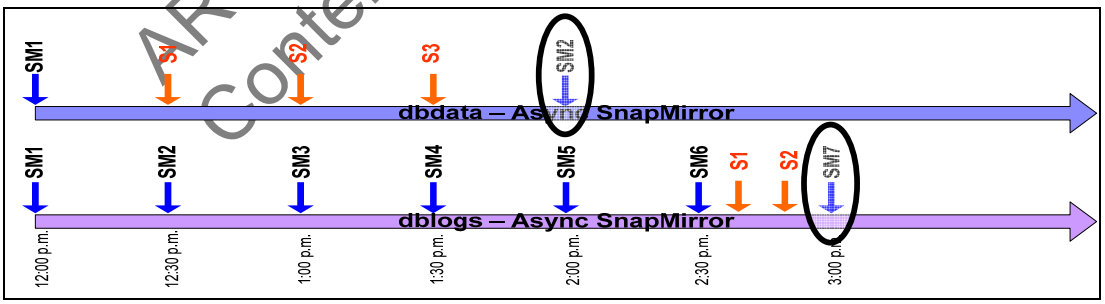

Figure 3) Recovery from a SnapMirror points.

In order to perform recovery from the SnapMirror update points for database data and transaction volumes, you would need to complete the following steps:

At the Disaster Recovery site, mount the volumes that were used as the SnapMirror destination for the database volumes on the primary site, to a database server by executing the following command:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
<NetApp storage System Name>:<volume name> <mount point>
```

For example, to mount a volume named `dbdata` on a mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
ntapdst:/vol/dblogs /mnt/dblogs
```

The control file that resides on the transaction logs volume will have the latest database information, so you would copy that control file from the `dblogs` volume to the `dbdata` volume by executing the following commands on the database server:

```
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control01.ctl
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
```

Start the database instance and recover the database by executing the following commands on the database server:

```
sqlplus / as sysdba
sql>startup mount
sql>recover database
sql>alter database open
```

After completing the above steps, the database recovery from the SnapMirror point of the database volumes will be complete.

2) Recovery from an application-coordinated Snapshot copy of the data volume

This type of recovery may become necessary in rare cases, such as when database recovery is not possible from the SnapMirror update point. In this type of situation, the database can be recovered using application-coordinated Snapshot copies. The RPO will be the time from the last successful SnapMirror update of the `dblogs` volume. The RTO will vary based on the transaction logs that need to be applied to recover the database.

For example, in our test scenario, the RPO will be 5 minutes. Data after the last SnapMirror update point of the `dblogs` volume will be lost.

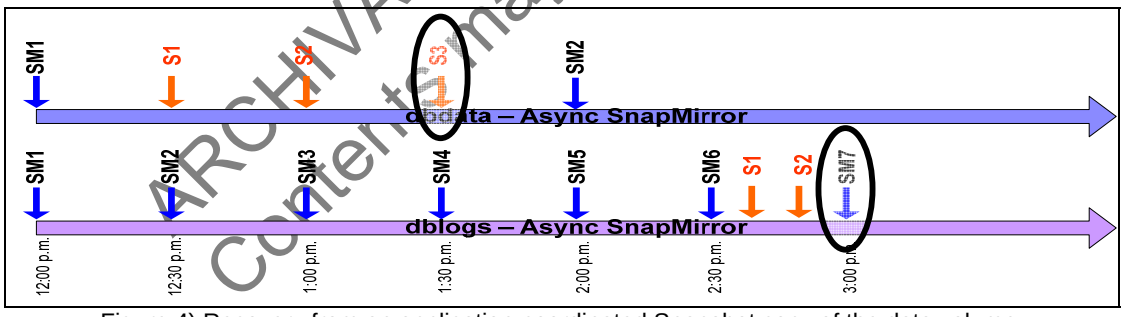

Figure 4) Recovery from an application-coordinated Snapshot copy of the data volume.

In order to complete database recovery from the application-coordinated Snapshot copy of the data volume and the SnapMirror update point of the transaction logs volume of the database, you would need to complete the following steps:

First you need to snap restore the volumes that hold data for the database from the application-coordinated Snapshot copies by executing the following command on the NetApp storage system:

`snap restore -s <snapshot name> <volume name>`

For example, to restore a volume named `dbdata` from an application-coordinated Snapshot copy named `s3`, you would execute the following command on the NetApp storage system:

```
snap restore -s s3 dbdata
```

Mount the NetApp storage system volumes that hold the database's data and transaction logs to a database host by executing the following command :

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768, tcp,vers=3,timeo=600
<netapp storage system name>:<volume name> <mount point>
```

For example, to mount a volume named `dbdata` on a mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
ntapdst:/vol/dbdata /mnt/dbdata
```

The control file that resides on the transaction logs volume will have the latest database information, so copy that control file from the `dblogs` volume to the `dbdata` volume by executing the following commands on the database server:

```
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control01.ctl
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
```

Start the database instance and perform database recovery by executing the following commands on the database server:

```
sqlplus / as sysdba
sql>startup mount
sql>recover database
sql>alter database open
```

After completing the above steps, the database recovery from an application-coordinated Snapshot copy of the database data volume and the SnapMirror point of the transaction logs volumes will be complete.

3) Recovery from an application-coordinated Snapshot copy of the transaction logs volume

This type of recovery becomes useful in some cases, such as when recovery is desired to a point in time before the SnapMirror update of the transaction logs volume, or when the recovery is not possible from the SnapMirror update point of the transaction logs volume. In such cases, the application-coordinated Snapshot copy can be used to recover the database. The RPO (Recovery Point Objective) is determined from the time the Snapshot copy was created to the time the disaster started and the RTO (Recovery Time Objective) will vary based on the transaction logs that need to be applied for the database recovery.

For example, in our Disaster Recovery scenario, the disaster stuck at 3:05 p.m. and an application-coordinated Snapshot copy named `S2'` from which we want to restore the `dblogs` volume was created at 2:50 p.m. therefore, the RPO will be15 minutes. The work done after the application-coordinated Snapshot point will be lost.

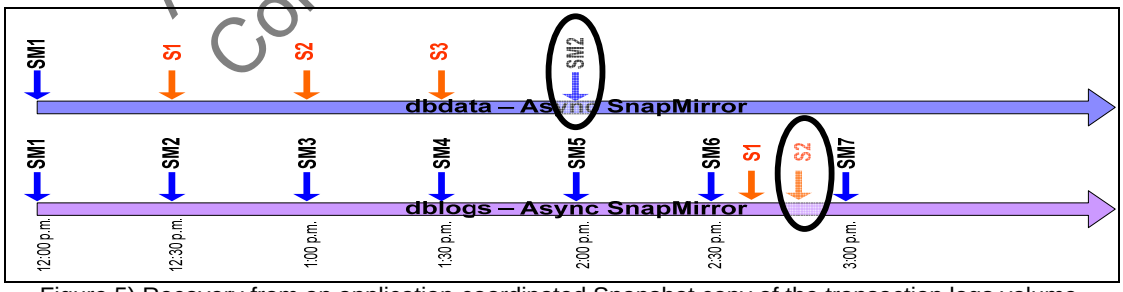

Figure 5) Recovery from an application-coordinated Snapshot copy of the transaction logs volume.

In order to perform database recovery using the SnapMirror point of the database's data volume and an application-coordinated Snapshot copy of the transaction logs volume, you would need to complete the following steps:

Restore the transaction logs volume from an application-coordinated Snapshot copy by executing the following statement on the NetApp storage system:

```
snap restore -s <snapshot name> <volume name>
```

For example, to restore a volume named `dblogs` from an application-coordinated Snapshot copy named `'s2'`, you would execute the following command on the NetApp storage system:

```
snap restore -s s2 dblogs
```

After restoring the database transaction logs volume, you need to mount the data and transaction logs volumes used for the database to a database host by executing the following command :

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
<netapp storage system name>:<volume name> <mount point name>
```

For example, to mount a volume named `dbdata` on a mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
ntapdst:/vol/dbdata /mnt/dbdata
```

The control file that resides on the transaction logs volume will have the latest database information, so copy that control file from the `dblogs` volume to the `dbdata` volume by executing the following commands on the database server:

```
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control01.ctl
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
```

Start the database instance and perform database recovery by executing the following commands on the database server:

```
sqlplus / as sysdba
sql>startup mount
sql>recover database
sql>alter database open
```

After completing the above steps, the database recovery from an application-coordinated Snapshot copy of the transaction logs volume and the SnapMirror point of the data volumes will be complete.

4) Recovery from application-coordinated Snapshot copies of the data and transaction logs volume

This case is variation of the previous option, "3). Recovery from an application-coordinated Snapshot copy of the transaction logs volumes." The only difference is that the data volume for the database is also restored from an application-coordinated Snapshot copy. The RPO is determined from the time that the Snapshot copy of the transaction logs volume was created to the time the disaster started. The RTO will vary based on the transaction logs that need to be applied for the database recovery.

For example, in our Disaster Recovery scenario, the disaster started at 3:05 p.m. and the application-coordinated Snapshot copy from which we want to restore the transaction logs volume was created at 2:50 p.m. Therefore, the RPO value will be15 minutes.

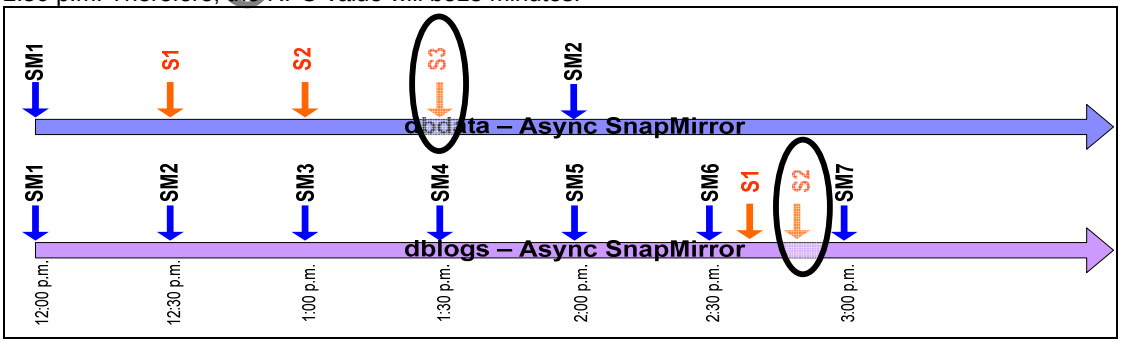

Figure 6) Recovery from the application-coordinated Snapshot copies of the `dbdata` and `dblogs` volumes.

In order to perform recovery from application-coordinated Snapshot copies of database data and transaction logs volumes, you would need to complete the following steps:

Restore the database data and transaction logs volume from application-coordinated Snapshot copies by executing the following command on the NetApp storage system:

snap restore -s <snapshot name> <volume name>

For example, to restore a volume named `dblogs` from an application-coordinated Snapshot copy named `'s2'`, you would execute the following command on the NetApp storage system:

```
snap restore -s s2 dblogs
```

Restore all the data as well as the transaction logs volumes using the above command.

After restoring the data and transaction logs volumes, you need to mount them to a database host by executing the following command:

mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600 <netapp storage system name>:<volume name> <mount point name>

For example, to mount a volume named `dbdata` on mount point `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
ntapdst:/vol/dbdata /mnt/dbdata
```

The control file that resides on the transaction logs volume will have the latest database information, so copy that control file from the `dblogs` volume to the `dbdata` volume by executing the following commands on the database server:

```
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control01.ctl
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
```

Start the database instance and perform database recovery by executing the following commands:

```
sqlplus / as sysdba
sql>startup mount
sql>recover database
sql>alter database open
```

After completing the above steps, the database recovery from an application-coordinated Snapshot copy of the data as well as the transaction logs volume will be complete.

12. DATABASE RECOVERY USING SNAPMIRROR ASYNC AND SYNC

The database Disaster Recovery strategy may include a mix of SnapMirror Async, Sync, and Semi-Sync modes of data replication. To simulate a Disaster Recovery strategy using mixed modes of SnapMirror, we configured the data volume replication using SnapMirror Async and the transaction logs volume replication using SnapMirror Sync mode. The SnapMirror update and application-coordinated Snapshot schedule has been summarized in following table.

Flexible Volume	SnapMirror	App. Coordinated Snapshot
Dbdata	**2 hrs**	**30 min**
Dblogs	**Sync**	**-**

To simulate a disaster at the primary site and perform recovery at the Disaster Recovery site that has replicated data using SnapMirror Async and Sync modes, you need to complete the following steps.

For this configuration, the entries in the `/etc/snapmirror.conf` file look somewhat similar to the following:

```
ntapsrc:dbdata ntapdst:dbdata – 0 0-22/2 * 1-5
ntapsrc:dblogs ntapdst:dblogs – sync
```

The argument 'sync' in the above entries controls whether a volume is Sync snapmirrored with its source or Async. If the argument is left out of the conf file, then the volume will be Async snapmirrored.

In order to simulate a disaster on the primary site and perform recovery on the Disaster Recovery site, you would need to complete the following steps:

Step 1 - Complete the SnapMirror base data transfer and run the load

Initialize the SnapMirror relationship for each volume used for the database and complete one-time baseline transfer by executing the following command on the NetApp storage system:

```
snapmirror initialize –S <source NetApp storage system>:<volume name>
<destination NetApp storage system>:<volume name>
```

For example, to initialize the baseline transfer for the volume named dbdata, you would execute the following command on the NetApp storage system:

```
snapmirror initialize –S ntapdst:dbdata ntapdst:dbdata
```

Monitor the SnapMirror status by executing the following command on the destination storage system:

```
snapmirror status
```

Let us assume our baseline transfer started at 12.00 p.m. and completed at 12.25 p.m.

Connect to the database and run the load by executing the following commands on the database server:

```
connect scott/tiger
execute scott.demobld(25000000);
```

Copy the do_snap, begin_bkup, and end_bkup scripts to a work directory on the database server and start taking coordinated Snapshot copies by executing the do_snap script on the database server.

```
./do_snap
```

The above command can be executed as a cron job from the database server.

Step 2- Simulate the disaster

Let us assume disaster strikes at 3:05 p.m. In order to simulate the disaster, we made the SnapMirror source NetApp storage system unavailable by executing the following command:

```
halt
```

At this point, the system state will look somewhat similar to the one shown in the Figure 7.

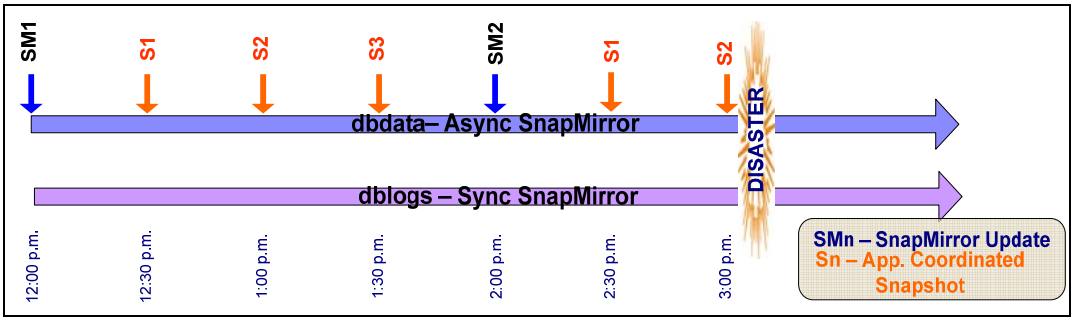

Figure 7) Database state at the time of disaster at the primary site.

After a disaster, the SnapMirror destination volumes corresponding to the database data will be in SnapMirror Snapshot consistent state. That means the data volumes at the Disaster Recovery site will have the same data, including application-coordinated Snapshot copies, as the corresponding source volume had at the time of successful SnapMirror update. The transaction logs volume is Sync snapmirrored; therefore, the corresponding volume at the Disaster Recovery site will have all the changes up to the point the disaster started. For our test scenario, the data volume named `dbdata` is Async snapmirrored and the transaction logs volume named `dblogs` is Sync snapmirrored; therefore, these two volumes will be in the following states:

`dbdata` – consistent with the source volume at the time of SnapMirror update at 2.00 p.m.
`dblogs` – consistent with the source volume at the time disaster started at 3.05 p.m.

Step 3 – Recovery

After the baseline transfer, the destination volumes are available to clients, but in a read-only state. The status of a destination will show that it is snapmirrored. A disaster makes the source unavailable; to use the destination volumes for writing as well as reading, you would need to end the SnapMirror relationship by executing the following command on the destination NetApp storage system:

```
snapmirror break <volume name>
```

For example, to make a volume named `dbdata` writable, you would execute the following command on the NetApp storage system:

```
snapmirror break dbdata
```

This command changes the destination volume's status from **snapmirrored** to **broken-off**, thus making it writable.

After a disaster, in this type of configuration, there are two possible options for the database recovery:

1) Recovery from the SnapMirror update point of the `data` volumes.
2) Recovery from an application-coordinated Snapshot copy of the data volumes.

1) Recovery from the SnapMirror update point of data volumes

The database's transaction logs volume is sync snapmirrored; as a result, all the changed data up to the point the disaster started is replicated to the destination volume. Therefore, after database recovery, there shouldn't be any data loss; the RPO (Recovery Point Objective) will be zero. The RTO (Recovery Time objective) will vary based on the transaction logs that need to be applied for recovery.

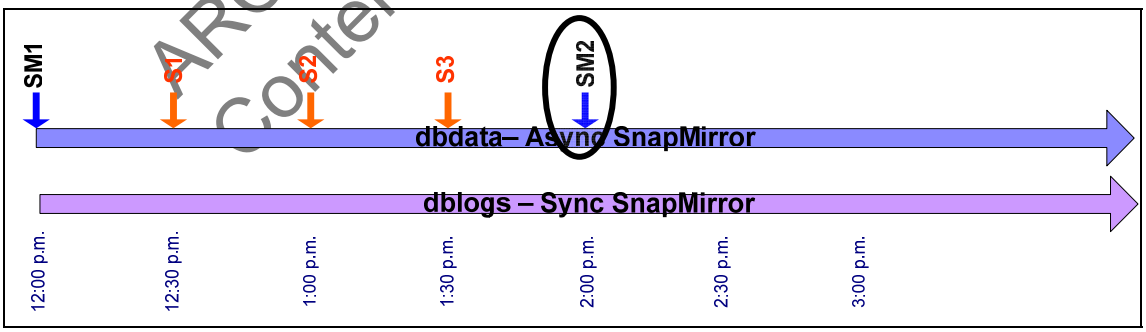

Figure 8) Recovery from the SnapMirror point of the `dbdata` volume.

To perform recovery from the SnapMirror update point of the volumes used for holding the database's data, you would complete the following steps on the database server:

You need to mount each volume used for the database to a database host by executing the following command :

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
<netapp storage system name>:<volume name> <mount point name>
```

For example, to mount a volume named `dbdata` that resides on a NetApp storage system named `ntapdst` and that needs to be mounted on a mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
ntapdst:/vol/dbdata /mnt/dbdata
```

The control file that resides on the transaction logs volume will have the latest database information, so copy that control file from the `dblogs` volume to the `dbdata` volume by executing the following commands on the database server:

```
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control01.ctl
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
```

Start the database instance and perform database recovery by executing the following commands on the database server:

```
sqlplus / as sysdba
sql>startup mount
sql>recover database
sql>alter database open
```

After completing the above steps, the database recovery will be complete.

2) Recovery from an application-coordinated Snapshot of the dbdata volume

This type of recovery may become necessary in rare cases, such as when database recovery is not possible from the SnapMirror update point. In this kind of situation, the database can be recovered using application-coordinated Snapshot copies of the data volume.

In this configuration, the database transaction logs volume is sync snapmirrored; as a result, all the changed data up to the point the disaster started is replicated to the destination volume. Therefore, after database recovery, there shouldn't be any data loss; RPO (Recovery Point Objective) will be zero. The RTO (Recovery Time objective) will vary based on the transaction logs that need to be applied for recovery.

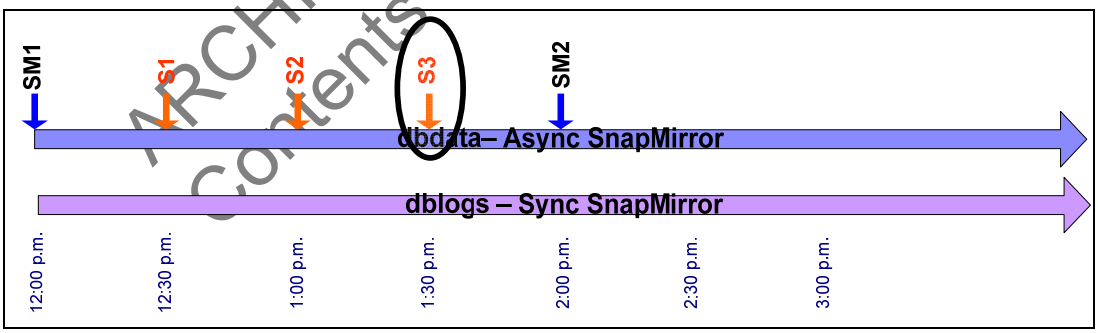

Figure 9) Recovery from an application-coordinated Snapshot copy of the `dbdata` volume.

In order to perform database recovery using an application-coordinated Snapshot copy of the volume holding the database's data, you would need to complete the following steps:

First you need to restore the volumes that hold data for the database from the application-coordinated Snapshot copy by executing the following command on the NetApp storage system:

```
snap restore -s <snapshot name> <volume name>
```

For example, to restore a volume named `dbdata` from an application-coordinated Snapshot copy named `s3`, you would execute the following command on the NetApp storage system:

```
snap restore -s s3 dbdata
```

Mount the NetApp storage system volumes that hold the database's data and transaction logs to a database host by executing the following command :

mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768, tcp,vers=3,timeo=600 <netapp storage system name>:<volume name> <mount point>

For example, to mount a volume named `dbdata` that resides on a NetApp storage system named `ntapdst` and that needs to be mounted on a mount point named `/mnt/dbdata`, you would execute the following command on the database server:

```
mount -o rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600
ntapdst:/vol/dbdata /mnt/dbdata
```

The control file that resides on the transaction logs volume will have the latest database information, so copy that control file from the `dblogs` volume to the `dbdata` volume by executing the following commands on the database server:

```
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control01.ctl
cp /mnt/dblogs/orcl/control03.ctl
/home/oracle/ora10g/dbdata/orcl/control02.ctl
```

Start the database instance and perform database recovery by executing the following commands on the database server:

```
sqlplus / as sysdba
sql>startup mount
sql>recover database
sql>alter database open
```

After completing the above steps, the database recovery will be complete.

13. CONCLUSION

NetApp SnapMirror is a proven data replication technology that offers great ROI. It provides flexible and robust methods to keep the recovery objectives controlled. The SnapMirror technology can be easily integrated with other NetApp technologies such as Cluster, Metro Cluster, and SyncMirror to further improve recovery objectives and ensure high availability.

APPENDIX: – A – SCRIPTS

Stored procedure script

```
CREATE OR REPLACE PROCEDURE demobld( rcount IN number DEFAULT 10) IS
--
-- Purpose: Generate database load by inserint specified number of rows in the
--          table
--
-- MODIFICATION HISTORY
-- Person      Date    Comments
-- ---------   ------  ------------------------------------------

BEGIN
   for i in 1..rcount  loop
      insert into scott.tab1 values (i,
'ABDFADFAFAFAFASFDASDFFFFFFFFFFFFFFFFFFFFFFFF');
      if mod(i,1000)=0 then
         commit;
      end if;
   end loop;
END; -- Procedure
/
```

Application-Coordinated Snapshot script (do_snap)

```
-------------------------------------------------------------------------------
--  Script Name: do_snap
--      Purpose: This script connects to the database and takes application –
--               Coordinated Snapshots
--        Usage:do_snap <NetApp storage system name> <vol name> <Snapshot name>
--
-------------------------------------------------------------------------------
if [ $# = 0 ]   then
   echo " USAGE: do_snap <NetApp storage system name> <vol name>  <snapshot name>"
   exit 1
fi
if [ $# = 1 ]  then
    echo "USAGE: do_snap <NetApp storage system name> <vol name> <snapshot name>"
    exit 1
fi
if [ $# = 2 ]  then
    echo " USAGE: do_snap <NetApp storage system name> <vol name> <snapshot name>"
    exit 1
fi
rsh $1 snap delete $2 $3
#-------------------------------------------------------------------------
# Put the database in hot backup mode
#-------------------------------------------------------------------------
sqlplus system/oracle @begin_bkup.sql
#-------------------------------------------------------------------------
# Create snapshot
#-------------------------------------------------------------------------
rsh $1 snap create $2 $3
#-------------------------------------------------------------------------
# Take the Tablespaces/datafiles out of hot backup mode
#-------------------------------------------------------------------------
sqlplus system/oracle @end_bkup.sql
rsh $1 snap list $2
echo ""
```

Application-Coordinated Snapshot script (begin_bkup)

```
--------------------------------------------------------------------------
-- Script Name: begin_bkup
-- Script Purpose: This script put the database in hot backup mode
--------------------------------------------------------------------------
alter database begin backup;
EXIT;
```

Application-Coordinated Snapshot script (end_bkup)

```
--------------------------------------------------------------------------
-- Script Name: end_bkup
-- Script Purpose: This script bring the database out of hot backup mode
--------------------------------------------------------------------------
alter database end backup;
EXIT;
```