



Oracle Backup and Recovery Best Practices for Dell SC Series Storage

Dell Storage Engineering
October 2015

Revisions

| Date | Description |
|-----------|--|
| 6/15/2011 | Initial draft |
| 4/12/2012 | Content and format change |
| 10/9/2015 | Added cloning of Oracle and Grid Infrastructure software |

Acknowledgements

Engineering: Mark Tomczik

THIS WHITE PAPER 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.

© 2011–2015 Dell Inc. All rights reserved. Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. 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 any proprietary interest in the marks and names of others.



Table of contents

| | | |
|---------|---|----|
| 1 | Introduction | 4 |
| 1.1 | Audience | 4 |
| 1.2 | Scope | 4 |
| 2 | Dell SC Series features | 5 |
| 2.1 | Data Instant Replay (DIR) | 5 |
| 2.2 | Consistency groups | 6 |
| 3 | Overview of Oracle storage | 7 |
| 3.1 | Oracle control files | 7 |
| 3.2 | Oracle online and archived redo logs | 7 |
| 3.3 | Database files | 8 |
| 3.4 | Oracle software | 8 |
| 4 | Overview of Oracle backups | 10 |
| 4.1 | Recovery Manager (RMAN) and user-managed backups | 10 |
| 4.2 | Physical and logical backups | 10 |
| 4.3 | Recovery Manager (RMAN) | 10 |
| 4.4 | User-managed backup | 11 |
| 5 | Oracle backup best practices with DIR | 13 |
| 5.1 | Using Oracle Recovery Manager (RMAN) | 13 |
| 5.1.1 | Using RMAN with a recovery catalog | 13 |
| 5.1.1.1 | Procedures and steps | 14 |
| 5.1.2 | Running RMAN backups with control files | 16 |
| 5.1.2.1 | Using the control files option | 17 |
| 5.2 | User-managed backups | 17 |
| 5.2.1 | Inconsistent or online backup | 18 |
| 5.2.1.1 | Procedures and steps | 19 |
| 5.2.2 | Consistent or cold backups | 21 |
| 5.2.2.1 | Cold database backup using Storage Center snapshots | 21 |
| 5.2.2.2 | Using snapshots to backup or clone RDBMS and Grid Infrastructure ORACLE_HOMEs | 24 |
| 5.3 | Snapshot interval and expiration times | 27 |
| 6 | Conclusion | 28 |
| A | Additional resources | 29 |
| A.1 | Technical support and resources | 29 |



1 Introduction

With Oracle databases surpassing terabytes in size, it is costly and time consuming to take backups of large mission-critical databases. Even if there is money to spend on multiple HBAs, backups may still require many hours to complete and can exhaust processors, memory, and disk resources on the enterprise servers that host the databases.

With traditional snapshot technologies, at least the same amount of storage used in the source servers has to be allocated up front for snapshot copies or clones. In some cases, more than twice the amount of storage for a copy or clone of the production database has to be allocated for the snapshot. In addition, there may be a limited number of snapshots that can be taken on a traditional SAN.

This white paper explains how to leverage Dell SC Series Data Instant Replay (DIR) with Oracle RMAN or user-managed backups to back up or offload backups onto another server, and to provision Oracle software throughout the enterprise by using software clones. In the process, backup windows and resource contention are reduced, and the need to pre-allocate additional space in advance is eliminated.

1.1 Audience

This paper is intended for database administrators, system administrators, and storage administrators that need to understand how to take advantage of SC Series features to reduce the demands on resources and to increase efficiencies within the Oracle environment. Readers should be familiar with Dell SC Series storage and have knowledge with and prior experience in configuring and operating the following:

- Oracle architecture
- Oracle 10g, 11g, and 12c
- Oracle Grid Infrastructure
- Oracle Recovery Manager (RMAN)
- Oracle user-managed backups
- General understanding of SAN technologies

1.2 Scope

This document explains Dell SC Series features, steps, how-tos, and best practices for combining Oracle with SC Series storage technology to meet and mitigate the challenges of backups and Oracle software cloning procedures.



2 Dell SC Series features

2.1 Data Instant Replay (DIR)

Data Instant Replay (DIR) is most often compared to snapshot technology and provides continuous, space-efficient data protection using a feature called Replays. Replays create point-in-time copies of volumes in which further changes to a volume are journaled, allowing the volume to be rolled back to its original state when the Replay was created. This not only saves disk space, but speeds local recovery of lost or deleted files. Replays can be mounted as volumes (called view volumes) for the sake of partial or full-volume data restore.

Policy-based schedules, with varying intervals and expiration, can be created to manage replays and provide greater recovery capabilities to a previous-known state, and there is no limit on the number or replays taken. When used in conjunction with the consistency groups feature, replays provide data integrity of data spanning multiple volumes.

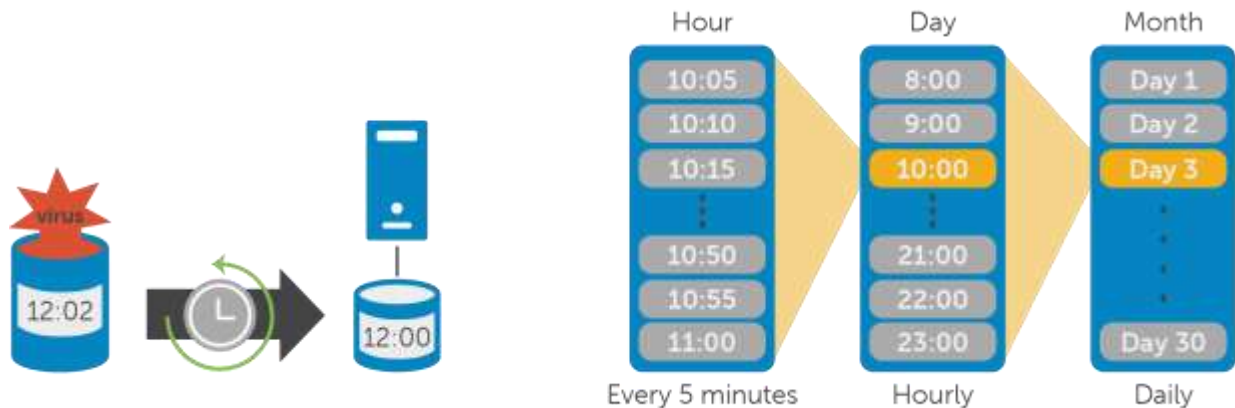


Figure 1 Data Instant Replay (DIR)

2.2 Consistency groups

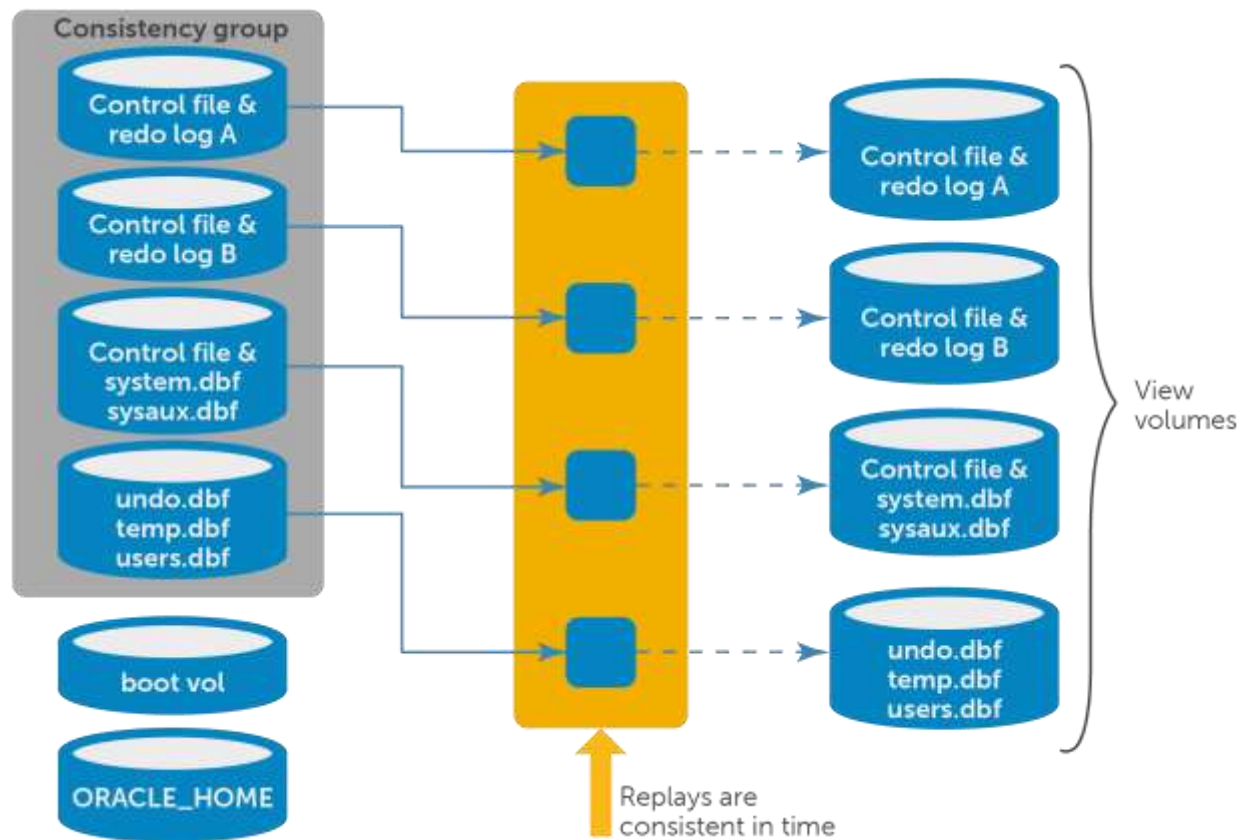


Figure 2 Consistency group and Oracle

Note: The consistency group feature allows storage administrators to take a snapshot of an Oracle database atomically. When creating a snapshot of an open database, you must ensure that all storage volumes (LUNs) that make up your database be atomically snapped because of multiplexed control files and redo log files. Remember that Oracle writes to multiplexed control files and redo log files concurrently, so without a consistency group you cannot create a consistent snapshot of a running database.

If a consistency group is not used, the database must be configured with all control files in one volume and all redo log files in the same volume or another volume, but they cannot be spread across volumes, whether file system or Oracle ASM is used. The consistency group feature provides the ability to create a usable snapshot of a database with control files and redo log files spread across volumes, which safeguards against having a single point of failure.

3 Overview of Oracle storage

Oracle requires the use of both physical and logical structures for storage. Physical structures are the files that make up the database and software. They consist of control files, redo logs, archive logs, datafiles, tempfiles, and other miscellaneous files, as well as Oracle binaries, libraries, and configuration files.

3.1 Oracle control files

The database control file, in relation to other Oracle database files, is a small Oracle proprietary-formatted binary file that contains critical information necessary for the database to operate successfully. If for some reason the control file is not accessible, the database cannot function properly.

A control file contains information about the associated database that is required for access by an instance, both at startup and during normal operation. Control files are updated continuously by Oracle during database use, so they must be available for writing whenever the database is open. Control file information can be modified only by Oracle; no database administrator or user can edit a control file.

Control files contain database information such as the database name and ID, the timestamp of database creation, tablespace information, redo log file information, and transaction sequence number. They also record information about checkpoints which are created every three seconds by the checkpoint process (CKPT). Checkpoints record positions in the redo log and are used during database recovery to instruct Oracle that all redo entries recorded before this point in the redo log group are not necessary for database recovery.

Oracle enables multiple, identical control files to be open concurrently and written for the same database. By storing multiple control files for a single database on different disks, the single point of failure with respect to losing a control file is mitigated.

3.2 Oracle online and archived redo logs

Another critical physical structure for recovery operations is the redo log file. Every instance of an Oracle database must have at least two redo logs to protect the database in case of failure. Each redo log is filled with redo records. Redo records, also called a redo entries, are made up of a group of change vectors, each of which is a description of a change made to a single block in the database. Redo records are buffered in a circular fashion in the redo log buffer of the SGA and are written to one of the redo log files by the log writer (LGWR) database background process. Whenever a transaction is committed, LGWR writes the transaction redo records from the redo log buffer of the SGA to a redo log file, and assigns a system change number (SCN) to identify the redo records for each committed transaction.

These redo log groups, multiplexed or not, are called an instance thread of redo. In typical configurations, only one database instance accesses an Oracle database, so only one thread is present. In an Oracle Real Application Cluster (RAC) environment, two or more instances concurrently access a single database and each instance has its own thread of redo. The two redo groups are requested to guarantee that one is always available for writing while the other is being archived (if ARCHIVELOG mode is enabled).



Oracle recommends that redo log files of an instance be multiplexed to safeguard against damage to any single file. When multiplexed, LGWR concurrently writes the same redo log information to multiple identical redo log files, thereby eliminating a single point of redo log failure.

When a redo log is filled up, the Oracle ARCH process is called to copy the redo log to an archive destination.

3.3 Database files

Database files contain all the data stored within the database, including all database metadata that describes the application data.

3.4 Oracle software

Oracle software reside in the ORACLE_HOME directories and are used to run and manage the Oracle environment. ORACLE_HOME directories are created as part of the preinstallation steps of installing Oracle and by running the runInstaller.sh script from the staging directory that contains the unzipped software downloads from Oracle.

Software staging directory:

```
# ls -ltr /u01/sw_oracle/11203
total 2442056
drwxr-xr-x. 8 oracle oinstall      4096 Sep 22  2011 database
-rw-r--r--. 1 root   root        1142195302 Feb 22  2012 p10404530_112030_Linux-
x86-64_2of7.zip
-rw-r--r--. 1 root   root        1358454646 Feb 22  2012 p10404530_112030_Linux-
x86-64_1of7.zip
# ls -ltr /u01/sw_oracle/11203/database
total 64
-rwxr-xr-x.  1 oracle oinstall  5466 Aug 23  2011 welcome.html
drwxr-xr-x. 12 oracle oinstall  4096 Sep 18  2011 doc
-rwxr-xr-x.  1 oracle oinstall  3226 Sep 22  2011 runInstaller
drwxr-xr-x.  2 oracle oinstall  4096 Sep 22  2011 rpm
drwxr-xr-x.  2 oracle oinstall  4096 Sep 22  2011 response
```



ORACLE_HOME directories for RDBMS and Grid Infrastructure:

```
# ls -ltr /u01/app/oracle/product/11.2.0/
```

```
total 8
```

```
drwxr-xr-x. 74 oracle oinstall 4096 Oct  1 13:21 dbhome_1
```

```
drwxr-xr-x. 73 oracle oinstall 4096 Oct  1 13:32 dbhome_2
```

```
# ls -ltr /u01/app/grid/product/11.2.0/
```

```
total 4
```

```
drwxr-x---. 67 grid oinstall 4096 Sep 25 19:24 grid
```



4 Overview of Oracle backups

Backups in Oracle can be classified as either logical or physical, and as a Recovery Manager (RMAN) or user-managed backup. This section discusses each of these types.

4.1 Recovery Manager (RMAN) and user-managed backups

Recovery Manager (RMAN) is an Oracle provided utility for backing-up, restoring, and recovering Oracle databases. Recovery Manager uses information about the database to automatically locate, back up, restore, and recover datafiles, control files, and archived redo logs. RMAN backups can be created using the RMAN command line interface or Oracle Enterprise Manager. A recovery catalog is an optional database schema that RMAN uses to store its repository data, which contains information about the databases RMAN will be managing, the RMAN backups performed, and the scripts needed by RMAN for backup, restore, and recovery operations. If a recovery catalog is not used, the controlfile of the target database is the sole source of metadata needed for backup, restore, and recovery operations.

User-managed backups consist of any strategy in which RMAN is not used as the principal backup, restore, and recovery tool. The basic user-managed backup strategy is to make periodic backups of datafiles, control files, redo, and archived logs with operating system commands. User-managed backups can also be used for backing up and restoring the Oracle software, and for cloning the ORACLE_HOME directories of the Oracle and grid infrastructure users.

4.2 Physical and logical backups

Physical backups can be created from Oracle Recovery Manager (RMAN) or from an operating system utility or command. Physical backups are backups of physical files such as database files, control files, redo logs, and archived redo logs (if ARCHIVELOG mode is set). They also exist in one of two formats: a proprietary Oracle format generated by RMAN (BACKUP) or an image copy. The image copy can be created by RMAN (COPY) or by an operating system utility or command (such as UNIX cp) and is a file-to-file copy.

Logical backups are created by reading data from a database object(s) and writing them to an OS file. In the more traditional method, a full database export is performed with the Oracle Export utility. SQL*Plus could also be used to select data from objects and write data to a file.

Timesaver: Use DIR to create a replay (snapshot) of your Oracle database or Grid Infrastructure software for cloning or backup and recovery purposes.

4.3 Recovery Manager (RMAN)

RMAN is the preferred solution by Oracle for database backup, restore, and recovery operations. It can perform the same types of backup and recovery available through user-managed methods more easily, provide a common interface for backup tasks across different host operating systems, and offer a number of backup techniques not available through user-managed methods.



Timesaver: Offload RMAN backups from the database host by using DIR to create and map snapshots or clones to a backup server. Once this is complete, invoke RMAN to back up the database on the backup server. The idea of minimizing backup windows on the database server is to have a dedicated server solely for backups, thus reducing the load on the database server. This requires a RMAN recovery catalog be configured.

The following file types are supported by RMAN:

- Data files
- Control files
- spfiles
- Archived redo logs
- Image copies of datafiles and control files
- Backup pieces that contain RMAN backups

The following files cannot be backed up by RMAN. They must be backed up using normal operating system utilities.

- External files
- Network configuration files
- Password files
- Any Oracle or grid software files

4.4 User-managed backup

User-managed backups can be categorized as logical or physical, and consistent or inconsistent. The type of backup made depends upon business needs.

A database is in a consistent state after it has been shut down with the SHUTDOWN options of NORMAL, IMMEDIATE, or TRANSACTIONAL. When a database has been shut down with any of these options, Oracle guarantees that all redo records have been applied and committed to the database. If a backup is made from a consistent database, the backup is considered to be a consistent backup and it can be mounted and opened without requiring any type of recovery.

When making a user-managed backup of ORACLE_HOME and GRID_HOME, Oracle states that all Oracle resources or services running out of the ORACLE_HOME or GRID_HOME are shut down prior to creating the user-managed backup, especially if the backup is to be used as a cloning source. This is to ensure that all file handles and locks to any file in ORACLE_HOME and GRID_HOME are released prior to making the user-managed backup, thereby making it a consistent backup of the software.

An inconsistent database backup is made when the database was not first shut down using one of the three shutdown options described previously. If an inconsistent database backup is restored, Oracle must perform recovery on the database before it can be opened. If a database is running in NOARCHIVELOG mode, Oracle states that inconsistent backups should not be taken. If a database is running in ARCHIVELOG mode, and archive redo logs are backed up, inconsistent backups can be a strong



component of a recovery strategy. Oracle also recommends that if changes are to be made against the database during the backup, the database needs to be placed in BACKUP mode while the backup is made.

In VLDB systems, full-consistent-physical backups are rarely or never taken because of time constraints. However, there are times when Oracle does recommend taking a cold backup, such as before and after an upgrade, after new tablespaces/datafiles are added, after a recovery, and after a resetlogs. For more information and recommendations regarding backup methods, refer to the *Oracle Database Backup and Recovery Users Guide*, which can be located on <https://docs.oracle.com/>.

As for user-managed backups, create a snapshot (Replay) of the volumes that make up the database (such as datafiles, redo logs, and archives) using a consistent Replay profile. Then create a view volume for each of the volume Replays, map the view volumes to the backup server, and use user-managed backup software to send the data to the desired backup media. The snapshots (Replays) can be left alone on the SAN after creation and can be expired anytime depending on the business requirement. Dell SC Series DIR technology provides the capability to recover any damaged or lost files quickly and easily.



5 Oracle backup best practices with DIR

5.1 Using Oracle Recovery Manager (RMAN)

There are two ways to use RMAN for backup:

- With a recovery catalog database
- With control files

5.1.1 Using RMAN with a recovery catalog

A recovery catalog database should be created on a separate server, preferably on the backup server itself. The steps in this section describe how to perform a backup and recovery with RMAN on a backup server.

The following assumptions are made:

- Both hosts (production and backup) are connected to the same SC Series system where the production Oracle volumes are kept
- An RMAN catalog is used
- The backup and production hosts have Oracle Net connectivity to the Recovery Catalog database
- The backup host has a dedicated volume for RMAN backup sets and has NFS export back to the production host

Figure 3 shows an example configuration of production and backup hosts. The production host has 3 volumes dedicated for the production Oracle database and a NFS mount point. The backup host has two volumes dedicated to the Recovery Catalog database and one volume for the RMAN backup sets.

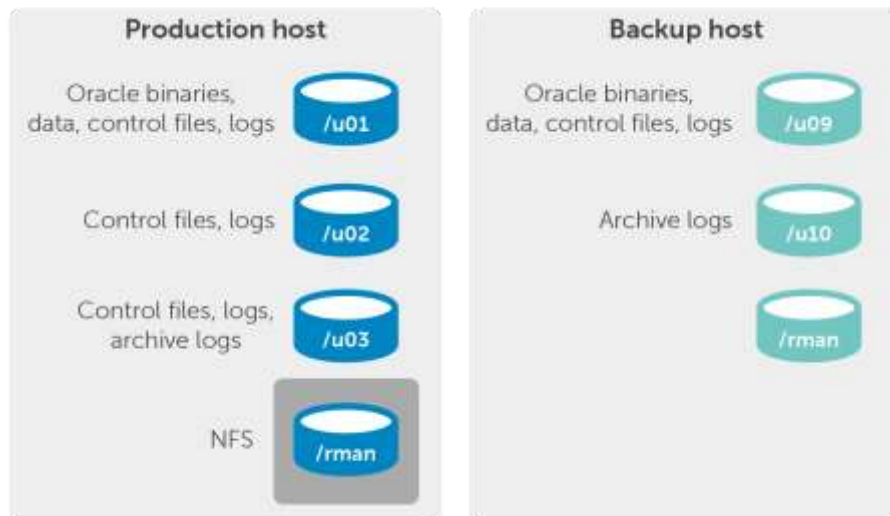


Figure 3 Example production backup configuration

5.1.1.1 Procedures and steps

This section illustrates two hosts:

Production host: This runs the production Oracle database and houses the production datafiles, online redo logs, control files, and archived redo logs.

Backup host: This is used to perform the RMAN backups and houses its own Oracle datafiles, online redo logs, control files, and archived redo logs for the RMAN catalog database.

On the **backup host**, follow these steps:

1. ssh to the production host, force the current redo log to be archived, and put the database in backup mode. This can be accomplished with a script.

See the following sample script named "startdb_backup.sh":

```
su - oracle << EOF
export ORACLE_SID=SID
export ORACLE_HOME=${ORACLE_HOME}
${ORACLE_HOME}/bin/sqlplus "/as sysdba" <<END
alter system archive log current;
alter database begin backup;
END
EOF
```

Execute the sample script "startdb_backup.sh":

```
ssh ${prod_host} /path/startdb_backup.sh
```

2. In Storage Center, create a consistency snapshot of the production database volumes.
3. ssh to the production database host, take the database out of backup mode, force a log switch, and archive the old redo log. This can be accomplished with a script.

See the following sample script named "enddb_backup.sh":

```
su - oracle << EOF
export ORACLE_SID=SID
export ORACLE_HOME=${ORACLE_HOME}
${ORACLE_HOME}/bin/sqlplus "/as sysdba" <<END
alter database end backup;
```



```
alter system archive log current;

END

EOF
```

Execute the sample script "startdb_backup.sh":

```
ssh ${prod_host} /path/endddb_backup.sh
```

4. Using the Storage Center GUI or CLI (CompCU), create view volumes from the snapshots created previously.
5. Using the Storage Center GUI or CLI (CompCU), map the view volumes to the backup host.
6. On the backup server, scan for the new devices, and mount the view volumes on the backup host with the same mount points as used on the production host. If running VxVM, make sure to import the Veritas disk groups. If running LVM, make sure to import the LVM disk groups. If running Linux EXT3 or EXT4, mount the volume in Linux. For ASM databases, make sure the `asm_diskstring` contains the proper strings of your ASM LUNs. For example, if using ASMLib on a database host, ASMLib needs to be installed on the backup host too so `asm_diskstring` does not need to be modified.
7. Start up the cloned production database, which resides on the view volumes, on the backup server with the MOUNT option:

```
SQL> startup mount;
```

8. Run RMAN backup to disk (/rman file system) on the backup server.

See the following sample RMAN script with a backup degree of parallelism 4:

```
rman catalog=rman/rman@catdb target=sys/oracle@SIDrun

{

configure device type disk parallelism 4;

allocate channel c1 type disk maxopenfiles 10 format
'/rman/PRODDB/%U.bak';

allocate channel c2 type disk maxopenfiles 10 format
'/rman/PRODDB/%U.bak';

allocate channel c3 type disk maxopenfiles 10 format
'/rman/PRODDB/%U.bak';

allocate channel c4 type disk maxopenfiles 10 format
'/rman/PRODDB/%U.bak';

backup as compressed backupset database plus archivelog;

}
```



9. When the backup completes, catalog all the backup sets created in the previous step and written to the directory `/rman` in the Catalog database.

See the following sample RMAN catalog script that catalogs 2 backup sets. This sample of code would need to be modified to back up all the backup sets created previously.

```
rman catalog=rman/rman@catdb target=sys/oracle@SID

run

{

catalog backuppiece "/rman/PRODDb/0m162of8_1_1.bak";

catalog backuppiece "/rman/PRODDb/0n162of9_1_1.bak";

}
```

10. When the catalog completes, shut down the production cloned database on the backup host. Then, unmount the file systems. In this example, unmount `/u01`, `/u02`, and `/u03`. If using VxVM or LVM, clean up your disk groups or volume groups before removing any volume mappings. If cleanup is not done properly, the next time a backup is created (create a new snapshot and corresponding view volume, and then map the view volume to your backup host), the operating system may be confused with previous configurations.
11. Use backup software to send all the backup sets from the `/rman` path to tape for offsite storage. Depending on the business backup policy, sending the daily backup tapes to offsite storage may not be allowed because the price of tapes can be expensive over time. Also, if a recovery is needed, a request may have to be made to acquire the offsite tape which could add extra time to the recovery process and may not be an acceptable recovery solution for the business. If the backup sets are left on the SC Series storage and Data Progression is used to migrate the backup sets to the lowest tier of storage (SATA drives), recovery can still be performed instantly since the files are readily available. RMAN retention can be configured as well, reducing the strains of space consumption.

5.1.2 Running RMAN backups with control files

The procedures for running RMAN backups with control files are similar to the procedures outlined in the section 5.1.1.1 except for step 9. Follow step 1 through step 8 inclusively using the above procedures. When step 9 is reached, run the catalog command directly on the production host. This will make the production database control files aware of the new backup sets.

The following assumptions are made:

- Both hosts (production and backup) are connected to the same Storage Center where the production Oracle volumes are kept.
- The backup host has a dedicated volume for RMAN backup sets and has NFS export back to the production host.



The following section provides an example of how to catalog the backup sets directly on the production host.

5.1.2.1 Using the control files option

See the following sample RMAN catalog script that catalogs two backup sets. This sample of code would need to be modified to back up all the backup sets created in section 5.1.1.1, step 8.

```
rman target /  
  
run  
  
{  
  catalog backuppiece "/rman_backups/PRODDb/0m162of8_1_1.bak";  
  catalog backuppiece "/rman_backups/PRODDb/0n162of9_1_1.bak";  
}
```

After the catalog command completes, run step 10 and 11 to complete the backup process.

5.2 User-managed backups

User-managed backups are candidates for databases that reside in Oracle Automatic Storage Management (ASM) or reside on file systems like ext4. Methods for creating the user-managed backups are varied and most likely will be different between an ASM and non-ASM resident.

User-managed backups can also be used for backing up ORACLE_HOME and GRID_HOME. When deploying and managing the Oracle and grid installations across an enterprise environment, user-managed backups of ORACLE_HOME and GRID_HOME can be used to greatly reduce the amount of time a database administrator (DBA) spends on and simplifies the process of managing the environment, especially when tens or even hundreds of database servers or clustered servers exist. Conversely, it could be argued that if adding one additional server to the environment, the traditional and interactive methods of using the Oracle Universal Installer (OUI) or the Provisioning Pack of Enterprise Manager should be used. However, there is a strong possibility that by using the manual cloning process, you will save additional time and gain efficiencies in ongoing support in the environment because the new server will have the identical Oracle footprint of the clone image. With this said, it should be noted that manual cloning is not a panacea for installing Oracle, and it might be more advantageous to use Oracle Enterprise Manager cloning. Determining which method to use will be governed by evaluating requirements.

In the following scenarios, cloning ORACLE_HOME of the RDBMS and Grid Infrastructure is beneficial:

- Prepare the environment once, and deploy it to many hosts simultaneously
- An xterm environment is not needed because there is no graphical user interface
- All patches can be applied in a single step, rather than going through the installation process multiple times for each patch
- Cloning guarantees a method to accurately repeat the same Oracle RDBMS and Clusterware installation on multiple server clusters



The process by which clones of ORACLE_GRID and ORACLE_HOME are made is driven by the requirements and limitations of the Oracle software, the physical and software architecture of the enterprise, Oracle versions, the use of RAC or non-RAC, file systems, logical volume managers, and other factors, but it is not driven by Dell Storage Center. It is imperative that the Oracle procedure for cloning Oracle and grid homes be reviewed and followed to ensure that the backup/clones are usable clones.

The following example shows how to use user-managed backup of a database, and then discusses using user-managed backups to clone a GRID_HOME and several ORACLE_HOME all residing in /u01.

The following assumptions are made:

- Both hosts (production and backup) are connected to the same Storage Center where the production Oracle volumes are kept.
- The source and target databases are 11gR2

Below is an example configuration of a production database server and backup server. The production server has three volumes dedicated for the production Oracle database. The backup server has no additional volumes except for the boot volume.

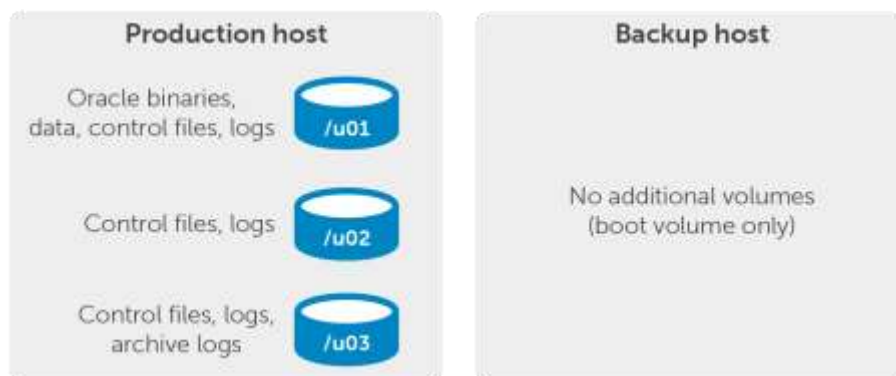


Figure 4 Example production user-managed backup configuration

5.2.1 Inconsistent or online backup

In order to create an online or hot backup, the database must run in archivelog mode. To determine the current archivelog mode, execute the following:

```
SQL> archive log list;
```

| | |
|------------------------------|---------------------|
| Database log mode | Archive Mode |
| Automatic archival | Enabled |
| Archive destination | /u02/oraarch/proddb |
| Oldest online log sequence | 35 |
| Next log sequence to archive | 36 |

Database log mode must be set to **Archive Mode** and **Automatic archival** must be set to **Enabled**. If they are not, execute the following commands on the source production database:

```
SQL> shutdown immediate;

SQL> startup mount exclusive;

SQL> alter database archivelog;

SQL> alter database open;

SQL> alter system set log_archive_start=TRUE [scope='spfile'];
```

5.2.1.1 Procedures and steps

Throughout this section, two hosts are illustrated:

Production host: This runs the production Oracle database and houses the production datafiles, online redo logs, control files, and archived redo logs.

Backup host: This is used to perform user-managed backups.

Perform the following backup procedure on the backup host:

1. ssh to the production host, force a log switch, and put the database in backup mode. This can be accomplished using a script.

See the following sample script named "startdb_backup.sh".

```
su - oracle << EOF

export ORACLE_SID=SID

export ORACLE_HOME=${ORACLE_HOME}

${ORACLE_HOME}/bin/sqlplus "/as sysdba" <<END

alter system archive log current;

alter database begin backup;

alter database backup controlfile to 'bkppath/control.ctl';

alter database backup controlfile to trace;

END

EOF
```

Execute the sample script, "startdb_backup.sh".

```
ssh ${prod_host} /path/startdb_backup.sh
```



2. In Storage Center, create a consistency snapshot of the production database volumes.
3. ssh to the production host, take the database out of backup mode, and force another log switch and archive the previous redo log. This can be accomplished using a script.

See the following sample script named "endddb_backup.sh":

```
su - oracle << EOF

export ORACLE_SID=SID

export ORACLE_HOME=${ORACLE_HOME}

${ORACLE_HOME}/bin/sqlplus "/as sysdba" <<END

alter database end backup;

alter system archive log current;

END

EOF
```

4. Using the Storage Center GUI or CLI (CompCU), create view volumes from the snapshots created previously.
5. In Storage Center, create a snapshot of the volumes that hold the archive redo logs of the production database.
6. Using the Storage Center GUI or CLI (CompCU), map the view volumes created from the snapshots to the backup host.
7. On the backup server, scan for the new devices, and mount the view volumes on the backup host with the same mount points as used on the production host. If running VxVM, then make sure to import the Veritas disk groups. If running LVM, then make sure to import the LVM disk groups. If running Linux EXT3, just do a mount.
8. Since the source volumes were dedicated to the Oracle database, use your backup software to send all the files from these three mount points to your backup media. If the list of database files to backup is needed, execute the following commands against the production source database in SQL*Plus:

```
select name from v$datafile

select name from v$tempfile

select name from v$controlfile

select member from v$logfile;
```

Archived redo can be found in the location returned by the following command:

```
select name
      , value
      from v$parameter
```



```
where lower(name) like '%archive%dest%'
and value is not null;
```

Other Oracle configuration files like the password file and network files should be backed up also. For more information and recommendations regarding backup methods, refer to the *Oracle Database Backup and Recovery Users Guide*, which can be located on <https://docs.oracle.com/>.

9. Depending on the business backup policy, sending the daily backup tapes to offsite storage may not be allowed because of the price of tapes can be expensive over time. Also, if a recovery is needed, a request to acquire the offsite tape needs to be made which will add extra time to the recovery process and may not be an acceptable recovery solution for the business. If the backup sets are left on the Storage Center and Data Progression is used to migrate the backup sets to the lowest tier of storage (SATA drives), recovery can still be performed instantly since the files are readily available. Make sure to set the snapshot expiration time appropriately so that only the required snapshots are retained.

5.2.2 Consistent or cold backups

Performing cold backups with Data Instant Replay is an easy process. The following example shows how a cold backup of an Oracle database is taken using DIR, and also an example of how a user-managed backup can be used for cloning ORACLE_HOME and GRID_HOMES.

5.2.2.1 Cold database backup using Storage Center snapshots

Throughout this section, two hosts are illustrated:

Production host: This runs the production Oracle database and houses the production datafiles, online redo logs, control files, and archived redo logs.

Backup host: This is used to perform user-managed backups.

Perform the following backup procedure on the **backup host**:

1. ssh to the production host, and shutdown the database

See the following sample script named "shutdown_db_backup.sh".

```
su - oracle << EOF
export ORACLE_SID=SID
export ORACLE_HOME=${ORACLE_HOME}
${ORACLE_HOME}/bin/sqlplus "/as sysdba" <<END
shutdown immediate;
END
EOF
```



Execute the sample script, "startdb_backup.sh":

```
ssh ${prod_host} /path/shutdown_db_backup.sh
```

2. In Storage Center, create a consistency snapshot of the production database volumes.
3. ssh to the production host, and start the database. This can be accomplished using a script.

See the following sample script, "start_db.sh":

```
su - oracle << EOF
export ORACLE_SID=SID
export ORACLE_HOME=${ORACLE_HOME}
${ORACLE_HOME}/bin/sqlplus "/as sysdba" <<END
startup;
END
EOF
```

Execute the sample script, "start_db.sh":

```
ssh ${prod_host} /path/start_db.sh
```

4. Using the Storage Center GUI or CLI (CompCU), create view volumes from the snapshots created above.
5. Using the Storage Center GUI or CLI (CompCU), map the view volumes created from the snapshots to the backup host.
6. On the backup server, scan for the new devices, and mount the view volumes on the backup host with the same mount points used on the production host. If running VxVM, make sure to import the Veritas disk groups. If running LVM, then make sure to import the LVM disk groups. If running Linux EXT3 or EXT4, just do a mount.

You should see something similar to Figure 5 after mapping the snapshots and mounting them.



Figure 5 Consistent cold backup configuration

7. Since the source volumes were dedicated to the Oracle database, use your backup software to send all the files from these three mount points to your backup media for offsite storage. If the list of database files to backup is needed, execute the following commands against the production source database in SQL*Plus:

```
select name from v$datafile
select name from v$tempfile
select name from v$controlfile
select member from v$logfile;
```

Archived redo can be found in the location returned by the following command:

```
select name
       , value
from v$parameter
where lower(name) like '%archive%dest%'
       and value is not null;
```

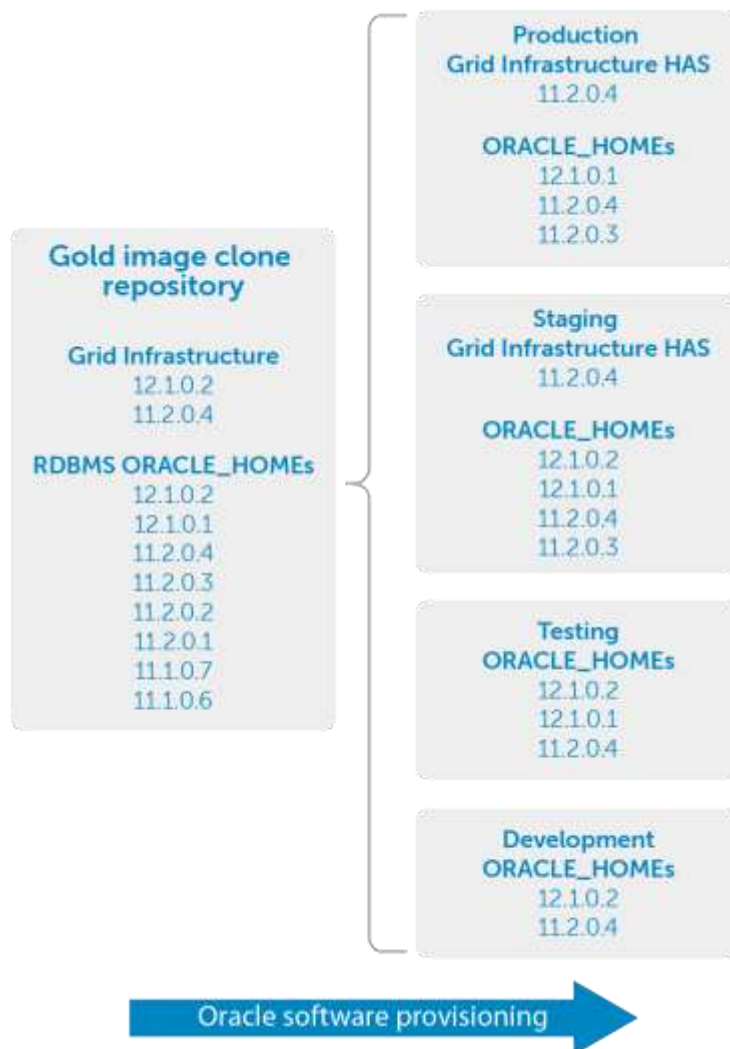
Other Oracle configuration files like the password file and network files should be backed up as well.

Depending on your business backup policy, you may not want to send the daily backup tapes to offsite storage because the price of tapes can be really expensive over time.



5.2.2.2 Using snapshots to backup or clone RDBMS and Grid Infrastructure ORACLE_HOMEs

Two database servers are used in this example, both running UEK6U4, 11.2.0.4 Grid Infrastructure standalone single instance, 11.2.0.4 RDBMS EE, and 11.2.0.3 RDBMS EE. All Oracle software and the OUI inventory reside in mount point /u01 on the source server. A snapshot of /u01 will be used to provision the Oracle and Grid Infrastructure software on a target server.



The following assumptions are made for the source and target servers:

1. Oracle Clusterware 11gR2 (11.2.0.4), RDBMS EE 11.2.0.4, RDBMS EE 11.2.0.3 is successfully installed on the source server
2. Grid Naming Service (GNS) is not used
3. Intelligent Platform Management Interface (IPMI) specification is not used

4. Clusterware and ASM spfile are stored in Oracle Automatic Storage Management (ASM)
5. No SCAN addresses are used
6. A single standalone instance is managed by the Oracle High Availability Service (HAS)
7. ASM Cluster File System (CFS) is not used
8. Cloning process is run by scripts in silent mode
9. User grid owns the Grid Infrastructure and software
10. User oracle owns the RDBMS software
11. The Linux user IDs of Oracle and Grid between the source and target servers are exactly the same
12. On the target server, users Oracle and grid belong to the same groups as they do on the source server
13. UEK6U4 x86-64 is installed on the target server exactly as it was on the source server
14. The OUI inventory location is the same between the source and target servers
15. The Oracle OS groups for DBA, OPER, and ASM are the same between the source and target servers.

Regardless of the configuration, it is very important that Oracle documentation is reviewed to make sure the cloning process is performed per Oracle recommendations and instructions.

Use the following backup and cloning procedure on the source server:

1. Log in to the source server as the user that owns the Grid Infrastructure.
2. Stop/shut down all Oracle and grid resources and services (including, but not limited to databases, listeners, applications, Oracle Clusterware, and ASM instance):
 - a. Use `srvctl` from the user (oracle) that owns the RDBMS software to shut down all non-ASM databases.

```
su - oracle

. oraenv

<SID>

srvctl status database -d <SID>

srvctl stop database -d <SID>

srvctl disable database -d <SID>
```

- b. Use `crsctl` from the user (grid) that owns the Grid Infrastructure to stop the rest of the components.

```
su - grid

crsctl disable has

crsctl stop has

crsctl status res -t
```



3. Using Storage Center, create a Replay of the volume that is mapped to /u01 on the source server. This Replay will be used as the clone source.
4. In Storage Center, create a view volume from the Replay created in the previous step and map the volume to the target server.
5. On the target server, scan for the new volume. This new volume contains the clone image of /u01 from the source server.
6. Update device mapper and /etc/fstab as necessary on the target server with the new volume.
7. On the target server, mount the new volume containing the clone image to mount point /u01.
8. Delete unnecessary files from clone image that pertain to the source server.
9. On the target server, update /etc/orainst.loc to point to the location of the OUI inventory.
10. Update listener.ora on the target server with the target server name.
11. Log in to the target server using the name of the user that owns the Grid Infrastructure on the source server and go to \$ORACLE_HOME/clone/bin directory.
12. Run the clone.pl script (the command below is on one command line):

```
perl clone.pl -silent ORACLE_BASE=/u01/app/grid
ORACLE_HOME=/u01/app/grid/product/11.2.0/grid
ORACLE_HOME_NAME=Orallg_gridinfrahome1
INVENTORY_LOCATION=/u01/app/oraInventory OSDBA_GROUP=asmdba
OSOPER_GROUP=asmoper OSASM_GROUP=asmadmin CRS=TRUE
```

13. Execute the script \$ORACLE_HOME/root.sh that resides in the \$ORACLE_HOME directory of the grid user.
14. Go to the directory \$ORACLE_HOME/bin and create the ASM database by running the following command:

```
./asmca -silent -configureASM -sysAsmPassword oracle -asmsnmpPassword
oracle -diskString 'ORCL:*' -diskGroupName CRS -diskList 'ORCL:CRS1' -
redundancy EXTERNAL -compatible.asm 11.2 -compatible.rdbms 11.2
```

15. Add the Grid oracle listener to HAS by executing the following from \$ORACLE_HOME/bin:

```
./srvctl add listener
./srvctl start listener
```

16. On the target server, log in to the user that owns that RDBMS software.
17. On the target server, run the following command for each of the \$ORACLE-HOMEs that reside in the clone image on /u01:

```
/u01/app/oracle/product/11.2.0/dbhome_1/root.sh
/u01/app/oracle/product/11.2.0/dbhome_2/root.sh
```

Depending on business needs and backup/cloning policies, a staging server may be used to stage all Oracle installations (Grid Infrastructure, RAC and standalone, and RDBMS EE) in /u01, then use a snapshot



of /u01 to provision Oracle to all the other servers in the enterprise. Doing this can greatly reduce the amount of time needed to manage Oracle software and guarantees that Oracle installations are identical across the entire enterprise.

5.3 Snapshot interval and expiration times

Depending on business needs, the size of the production database, the rate of change, and storage capacity, the information in Table 1 can be used as a starting point for snapshot intervals and expiration times.

Table 1 Sample snapshot schedule

| Snapshot interval | Snapshot expiration time |
|-------------------|--------------------------|
| Hourly | Daily |
| Daily | Weekly |
| Weekly | Monthly |
| Monthly | Yearly |



6 Conclusion

The need to reduce recovery time and improve backup speed will play a very important role in any IT organization as databases continue to grow and become strategically involved in business operations. Using Dell SC Series Data Instant Replay with RMAN and user-managed backups reduces the backup window and offloads expensive backup operations from the production host while conserving valuable CPU, memory, and disk resources. It also provides a mechanism to ensure that all Oracle installations are identical and provides a method to quickly and efficiently provision Oracle software on servers.



A Additional resources

A.1 Technical support and resources

For Copilot support of Dell SC Series products:

- SC Series [Customer Portal](#)
- Email: support@compellent.com (non-emergency business hours)
- Phone: 866-EZ-STORE (866-397-8673) (United States only)
- [Global online support](#)

The Dell SC Series [Portal](#) is an online portal for existing customers. A valid portal account is required to access the Knowledge Center. Once login to the portal, go to “Knowledge center”.

[Dell TechCenter](#) is an online technical community for IT professionals and is a great resource to discover and learn about a wide range of technologies such as storage, servers, networking, software, and cloud management.

