

ORACLE RMAN DESIGN BEST PRACTICES WITH EMC DATA DOMAIN

Abstract

Many organizations rely on Oracle databases as the foundation for their mission-critical applications and the ability to protect and recover this data or risk interruption to the business. EMC® Data Domain® deduplication storage integrates seamlessly with Oracle Recovery Manager (RMAN), enabling cost-effective data protection and disaster recovery solutions. This white paper provides technical considerations and specific configuration guidance for successfully deploying Data Domain systems as a direct target for Oracle RMAN.

November 2010



Copyright © 2010 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is”. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Part Number h8110

Table of Contents

Executive summary	5
Audience.....	5
Data Domain product background	5
Data Domain and Oracle.....	6
Benefits of using a Data Domain system as a target for Oracle RMAN	7
Deployment options	8
Oracle RMAN concepts and terminology	8
Types of backups	8
Backup set	9
RMAN parameters.....	9
MAXOPENFILES.....	9
FILESERSET.....	10
MAXPIECESIZE.....	10
Compression	10
Encryption	11
Channels and parallelism.....	11
ASM	12
General best practices for Oracle databases	13
Run in archivelog mode.....	13
Perform full Oracle backups on a regular basis	13
Keep archive logs and/or incremental backups in between each full backup.....	13
Keep one or more copies of the data on disk in separate locations.....	13
Use RMAN to help organize the various backups and simplify possible recovery	14
Summary of tuning parameter best practices for Oracle RMAN	14
Best practices for Oracle RMAN with Data Domain deduplication storage	15
Capacity	15
Networking.....	16
Security.....	17
Optimizing deduplication	17
Best practices for configuring RMAN options	18
Selecting the type of backup: Full, incremental, image	18
Logical directory structure for storing backups	19
FILESERSET.....	19
MAXPIECESIZE.....	19
Parallelism and channels	20
Compression	20
Encryption	21
Using multiple archivelog destinations.....	21

Very large database considerations.....	21
Conclusion.....	23
References	24

Executive summary

EMC® Data Domain® deduplication storage systems are designed and optimized specifically for backup and archive data.

There are many product attributes that enable Data Domain systems to excel at backup and archiving, including:

- Support for any conventional backup or archive application through generalized support for network-attached storage (NAS) interfaces over Ethernet, a virtual tape library (VTL) interface option over Fibre Channel, and product-specific interfaces such as NetBackup OpenStorage and EMC Data Domain Boost
- High-speed, inline deduplication using small, variable-sized sequences to identify and eliminate redundant data segments before storing to disk
- Integrated data protection technologies such as RAID 6, post-backup data verification, and periodic validation checks of existing data sets
- Automated replication of backup data for disaster recovery (DR) using cost-effective, low-bandwidth WAN links, which enables faster “time-to-DR” readiness

Data Domain systems are tuned for applications that perform sequential I/O, such as backups. As noted above, multiple interfaces are supported for these applications, but only the NAS interfaces are supported for use directly with Oracle Recovery Manager (RMAN). While many of the RMAN parameter recommendations will be applicable, environments requiring the use of the VTL and/or DD Boost interfaces will require RMAN integration with a supported data protection application, which is outside the scope of this paper.

The information in this paper can be applied to all supported versions of Oracle (releases 9, 10, or 11). Variances to the recommendations for a specific version will be noted where applicable.

Audience

This white paper provides technical information on the integration of Oracle RMAN with Data Domain deduplication storage systems. The reader should have a basic knowledge of Oracle management practices as well as familiarity with backup and recovery techniques using Oracle RMAN. A working knowledge of UNIX/Linux and/or Microsoft Windows and the networking components of each, along with a basic understanding of the setup and management of a Data Domain system, is also required. Theory and background will be presented where required for context or to further explain the best practice recommendations provided.

Data Domain product background

Data Domain deduplication storage systems have a number of unique capabilities that are designed to directly address the challenges of using disk for data protection and disaster recovery. Data Domain inline deduplication breaks the incoming data

stream into variable-length segments and uniquely identifies each one, and then compares the segments to previously stored data. If the segment is unique, it is compressed and stored on disk along with associated metadata. If an incoming data segment is a duplicate of what has already been stored, only the metadata reference to the existing segment is stored. The EMC Data Domain Data Invulnerability Architecture offers advanced data verification and data integrity, including RAID 6 protection, continuous fault detection, healing, and write verification to ensure maximum data integrity, availability, and recoverability. Finally, EMC Data Domain Replicator software transfers only the deduplicated and compressed changes across any IP network, requiring a tiny fraction of the bandwidth, time, and cost compared to traditional replication methods, enabling cost-effective disaster recovery.

Refer to the [Resources](#) section on page 24 for links to more detailed information on Data Domain technology, including:

- Data Domain SISL™ Scalability Architecture
- Data Domain Replicator software
- Data Invulnerability Architecture

On a Windows network, the Data Domain system presents shares via a Microsoft Common Internet File System (CIFS) protocol. On a UNIX or Linux network, the Data Domain system presents shares accessible via a Network File System (NFS) protocol. A single Data Domain system can present shares via all protocols simultaneously.

Data Domain and Oracle

Oracle RMAN is a built-in tool that allows the database administrator (DBA) to easily back up and recover data in an Oracle database. RMAN handles the coordination required to ensure that transaction integrity is preserved, and sufficient information is maintained to recover the database to any appropriate point. RMAN can create backup sets that comprise as much or as little recovery information as the DBA requires but usually include information from the database datafiles, control files, and redo and archived log files.

RMAN supports performing backups to a local tape drive¹ a local disk, or a NAS device, as well as integration with traditional enterprise backup applications, as shown in Figure 1.

¹ Support for DEVICE TYPE TAPE is provided by vendor-supplied RMAN plug-ins for various enterprise backup applications such as EMC NetWorker®, Oracle Secure Backup, Symantec NetBackup, IBM TSM, and others.

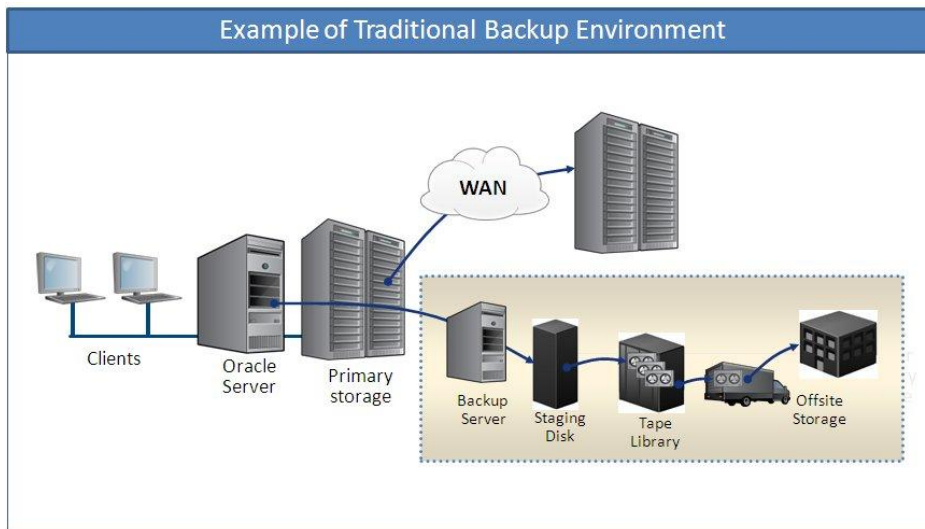


Figure 1. Example of a traditional backup environment

This paper will specifically focus on tuning the parameters to optimize backup performance and compression when using a Data Domain deduplication storage system as a direct NAS target for RMAN backups, as shown in Figure 2.

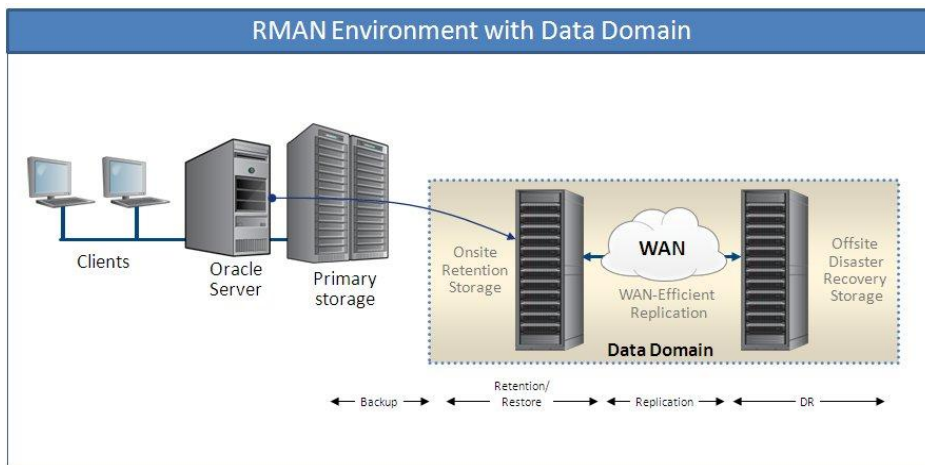


Figure 2. RMAN environment with Data Domain

Benefits of using a Data Domain system as a target for Oracle RMAN

By eliminating redundant data segments inline, Data Domain systems allow many more backups to be retained than would be possible using traditional storage. In particular, Data Domain systems use a variable-length segmentation process that is extremely efficient at finding identical segments within backups of monolithic files, such as Oracle datafiles.

The ability of the Data Domain system to store several weeks or months of full Oracle backups enables the DBA to implement a backup and recovery scheme with great flexibility and protection while consuming a minimal amount of physical storage. The integration of the Data Domain system into an Oracle/RMAN environment is seamless

since the Data Domain system presents itself either as an NFS or CIFS shared storage server. Oracle/RMAN already supports and documents this type of installation for effective RMAN storage.

If an enterprise backup software solution such as Oracle Secure Backup or EMC NetWorker is already in use, the Data Domain system can be seamlessly integrated into this environment as well. In this case, the Data Domain system can appear as an SBT_Tape device or a disk device to the enterprise backup software solution.

For critical Oracle environments, it is a best practice to replicate the production Oracle data to a secondary recovery location. The DBA has many options to choose from, including technologies from Oracle such as Oracle Data Guard, solutions offered by primary storage providers, and third-party solutions. Data Domain Replicator software offers extremely bandwidth-efficient replication that is also easy to deploy, enabling DBAs to leverage RMAN to provide disaster recovery capabilities for Oracle databases.

The primary benefit of Data Domain Replicator is the fact that only deduplicated and compressed data is transferred across the network. Because deduplication is happening inline, replication takes place while the RMAN backup process is still active. As the RMAN backup process proceeds, the unique segments and metadata representing each file in the backup set are queued for replication to the remote site, allowing the overall “time-to-DR” to be minimized. In many cases, replication is completed within a short period of time after the initial backup completes.

Deployment options

A Data Domain system can be deployed as a NAS device or as a VTL. Additionally, OpenStorage integration is available for environments using Symantec NetBackup 6.5 or a newer release. However, since RMAN does not have the capability to talk directly to a VTL or OpenStorage interface on the Data Domain system, it must use an enterprise backup application to provide that integration, if either of these protocols are desired.

For more detailed information on Data Domain technology, please refer to the list of supporting documents at the end of this paper.

Oracle RMAN concepts and terminology

Oracle RMAN has a set of important definitions and parameters that are unique and must be understood in order to fully utilize the application.

Types of backups

RMAN supports several kinds of backups.

- A **full backup**, as its name implies, is a backup of the entire database.

- A **partial backup** can be of a single or multiple portions of the database but not necessarily everything.
- An **incremental backup** uses a facility within Oracle to track changes to the database that occur between backup operations and will back up only those specific pieces of the database that have been altered.
- The final type of backup is an **image backup**, which makes precise copies of the Oracle database files. The result looks just like a copy of the files when viewed from the operating system. These generally take more space than a full backup because RMAN does not eliminate “whitespace” during this operation, but once deduplicated and compressed by a Data Domain system, the space used is similar to a full backup.

In each of these backups, RMAN processes the data to be backed up and consolidates it into one or more files called a backup set.

Backup set

A backup set is a set of one or more files that are written by RMAN as the output of a backup operation. The DBA can characterize a backup set in a number of different ways through different naming facilities as well as a separate tagging facility. It is important to note that a backup set can hold interleaved data from a number of different Oracle datafiles, which will have an impact on the ability of deduplication to identify redundancy.

RMAN parameters

There are many important options that affect how RMAN behaves when performing backup operations. Some of the most important are outlined here because they affect how RMAN will interact with a Data Domain deduplication storage system. For a more complete description of these options, refer to Oracle’s documentation and to Table 1 on page 14 for a summary of best practices.

MAXOPENFILES

This parameter controls the number of Oracle datafiles that RMAN can have open for reading at any one time.

An important effect of this parameter is on the amount of system memory consumed by RMAN during a backup operation. Reading from multiple files can improve throughput of the backup operation but will consume more resources.

To minimize the number of I/O buffers that RMAN allocates, it is important to adjust MAXOPENFILES to the smallest value possible that still achieves good performance.

MAXOPENFILES [=] integer	Controls the maximum number of input files that a BACKUP command can have open at any given time (the default is 8). Use this parameter to prevent “Too many open files” error messages when backing up a large number of files into a single backupset.
-----------------------------	--

FILESPERSET

This option, a parameter on the “begin backup” RMAN command, tells RMAN how many Oracle datafiles it can combine into a single backup file. In other words, it controls the multiplexing of Oracle data into the backup sets.

FILESPERSET [=] integer	When used with commands that create backupsets, this specifies the maximum number of files to include in each created backupset. By default, RMAN divides files among backupsets in order to make optimal use of channel resources. The number of files to be backed up is divided by the number of channels. If the result is less than 64, then it is the number of files placed in each backupset. Otherwise, 64 files will be placed in each backupset.
----------------------------	---

MAXPIECESIZE

This parameter defines the largest size a single file within the backup set can reach. It is a parameter on either the CONFIGURE CHANNEL or ALLOCATE CHANNEL commands.

This setting is important as it will have an impact on how replication

between Data Domain systems works and guidance is given in the [General best practices for Oracle databases](#) section.

MAXPIECESIZE [=] integer	Specifies the maximum size of each backup piece created on the defined channel. The size is specified in bytes, kilobytes(K), megabytes(M), or gigabytes(G). For example, if you set MAXPIECESIZE = 5000, RMAN sets the maximum piece size to 4 kilobytes, which is the lower four kilobyte boundary of 5000 bytes.
-----------------------------	---

Compression

RMAN provides two types of compression. All supported versions of RMAN can apply a binary compression algorithm (BZIP2) to the backup set. This results in less disk space being used at the cost of significantly greater CPU consumption during the backup operation. When using a Data Domain system as the target, RMAN lossless compression *should not* be used, as pre-compression of the backup streams will randomize the data patterns and defeat deduplication.

A second type of compression was first introduced in Oracle 10g. If RMAN encounters a block in a datafile being backed up that is not in use, it will not include it in the backup stream unless a certain set of conditions exists (see Oracle RMAN documentation for details). This feature is on by default and will not conflict with the ability of the Data Domain system to identify redundancy in the databases being backed up.

Recent versions of Oracle (11g) add additional compression options to RMAN. As with BZIP2, these are *not recommended* when backing up to a Data Domain system. Furthermore, there are forms of compression internal to the Oracle database (not involved with RMAN) that have no effect on how RMAN is run nor significantly affect deduplication.

Encryption

RMAN has the ability to encrypt backup sets to increase security. Any backup set can be encrypted except an “image” copy. Automatic encryption requires the configuration and setup of key management functions such as those provided by the Oracle Encryption Wallet. As with compression, encryption will randomize the data patterns and defeat deduplication, and *should not* be used when using a Data Domain system as the primary target for RMAN backups. If making subsequent copies of the backups to tape using a backup application, encryption can be applied to the data during the copy process.

Channels and parallelism

RMAN backs up Oracle datafiles using a series of separate processes, called channels, which run in parallel, that is, they are separate running programs and/or threads. Typically, while one process is waiting on I/O to complete for a given set of reads/writes, another process can be performing similar tasks against a separate set of files. By keeping these channels running simultaneously against data stored on different disks, the overall RMAN backup can complete in less time than if the entire process ran sequentially against a single file at a time.

Figure 3 shows an example that explicitly parallelizes a backup operation by specifying to Oracle Enterprise Manager (OEM) that four channels should be used.

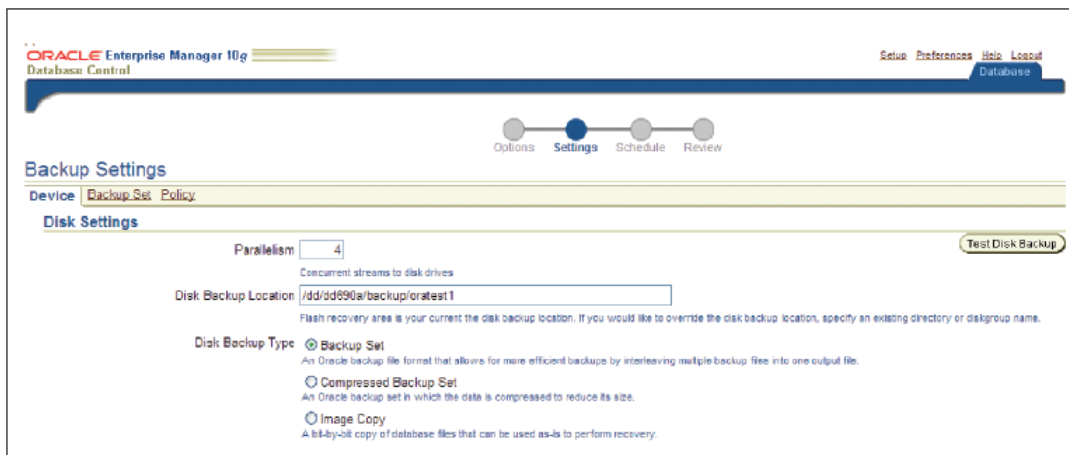


Figure 3. Parallelizing a backup operation

By specifying to OEM that Parallelism should be 4 and (on a different screen) that FILESPERSET should be set to 1, the result is a script like the one shown in Figure 4.

The DBA can choose to define a global setting for parallelism and allow RMAN to balance the I/O load as best it can, or the DBA can manually define these channels. The latter may allow for better performance, as well as other administrative effects, as specific knowledge of the storage environment can be taken into account.

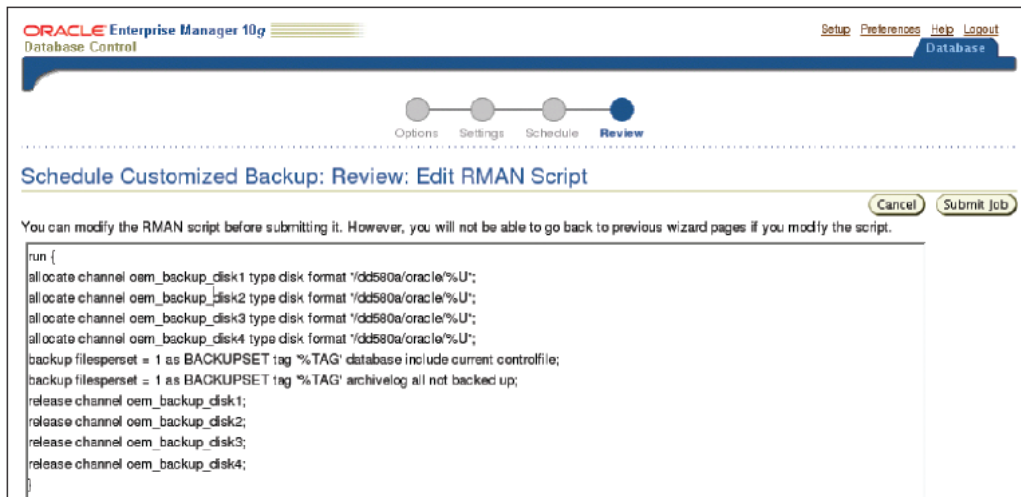


Figure 4. OEM script

Explicit definitions of channels also provide the DBA with a way to balance backup load across several separate network paths without employing network trunking or similar technologies. With separate mount points (NFS) or shares (CIFS) specified and specifying separate channels on each, the RMAN workload can be easily distributed across multiple network connections.

An example of multiple channels being used by RMAN to back up a database to different storage targets is shown in Figure 5.

The parameters just discussed can have a significant effect on Oracle RMAN performance when backing up to a Data Domain system. This includes the use of the Data Domain system as a direct target for RMAN (either using an NFS mount or a CIFS share) or if the files created are subsequently copied to the Data Domain system using a backup application or operating system utilities (for example, “cp” or “dd” for UNIX and Linux or “copy” for Windows). These effects, and combinations thereof, will be discussed in [Best practices for configuring RMAN options](#).

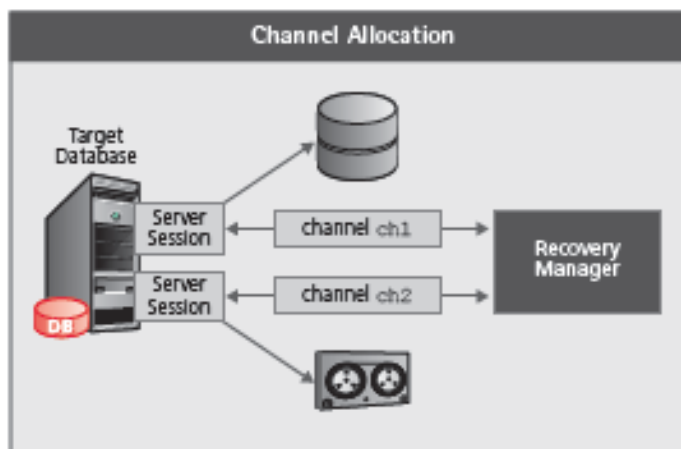


Figure 5. Channel allocation

ASM

Oracle’s Automatic Storage Management (ASM) was first introduced in Oracle 10g and was further enhanced and extended in 11g. ASM manages I/O to the Oracle datafiles, and while its use is outside the scope of this paper, it can have a bearing on performance in environments where it is used.

Also, since ASM is a database unto itself, it is important for a DBA to keep in mind the need to maintain effective backup and protection of the information required by ASM to provide its service to the rest of the Oracle environment.

General best practices for Oracle databases

Oracle environments range from small, experimental installations to large production environments that run multi-billion-dollar corporations. Because of this large difference in scale, a single recommendation is seldom applicable across all environments. However, some best practices for maintaining a reliable Oracle database exist, and are presented here for reference.

Run in archivelog mode

In this mode, Oracle does not delete the redo logfiles that track changes to the database. Instead it “archives” these files to secondary (and sometimes tertiary) locations when Oracle performs its periodic “logswitch”. Without this history of archived logfiles, the ability to recover an Oracle database to a point prior to an error is significantly constrained.

When provisioning primary storage, it is important to follow Oracle and storage vendor documentation for optimal storage of Oracle database files.

Perform full Oracle backups on a regular basis

The availability of multiple recovery options is important to the Oracle DBA as the size and importance of an Oracle database grow. It is important that all DBAs develop a practice for doing regular backups, the base of which is a full backup. Periodically practicing recoveries is also recommended to ensure that all of the pieces, processes, and skills are ready and available for the inevitable day when a critical recovery is needed. The various types and levels of Oracle backups will be covered in the section [Selecting the type of backup: Full, incremental, image](#).

Keep archive logs and/or incremental backups in between each full backup

It is critical to keep a history of archive logs in order to maximize recovery opportunities. A backup set consisting of a full database backup along with all the subsequent archive logs allows the DBA to recover the database to any point desired. Incremental backups also provide in-between recovery opportunities.

Keep one or more copies of the data on disk in separate locations

If only a single set of backup files is kept, there is a significant risk that something could happen to make it unavailable or unusable. To reduce that risk substantially, it is highly recommended that the backup practice arrange to keep multiple copies of the backup files in various locations.

Use RMAN to help organize the various backups and simplify possible recovery

To simplify all of these different operations, copies, and locations of files, the use of RMAN is highly recommended. It assists the DBA in keeping track of the state of all these various backup images as well as their locations. It also assists in making the decision on whether that particular backup information has become obsolete and can therefore be safely retired. Finally, it makes recovery for the DBA much simpler since RMAN keeps track of all of the files required for particular recovery or cloning operations.

Summary of tuning parameter best practices for Oracle RMAN

Here is a summary of the suggested best practice parameter values as detailed on previous pages.

Table 1. Summary of best practice settings

PARAMETER OR OPTION	SETTING
Type of backup	full
“FILESERSET”	1
“MAXPIECESIZE”	Depends on WAN replication requirements
Parallelism or Number of Channels	As many as the system can bear. Usually ranges between 3 and 5
Encryption	No
Compression (“AS COMPRESSED” clause)	No
NFS MOUNT OPTIONS	SETTINGS*
Sun Solaris	sw, hard, rsize=32768,wsize=32768,llock
Linux and other UNIX	sw, hard, rsize=32768,wsize=32768,nolock
ORACLE RAC environment	Add the “noac” or “actimeo=0” option to appropriate settings above (see particular Oracle release notes)
Access to mount point	Restrict access to a mount point to only the UNIX machines actually running RMAN and/or any standby, recovery systems
Replication	Use the Data Domain system to replicate backup sets to the remote DR site

MISCELLANEOUS OPTIONS	SETTINGS
Use of multiple Data Domain systems to back up very large databases	Be sure to carefully configure channels to ensure the same datafiles are always backed up to the same Data Domain system.
Multiplex archivelogs to the Data Domain system	Multiple copies of archivelogs provide added protection. Specifying the Data Domain system as an optional destination increases protection and gets the logs to the system even sooner than waiting until RMAN runs.

*The recommended rsize/wsize setting is a minimum. Larger values may be tried and have been seen to have beneficial results

Best practices for Oracle RMAN with Data Domain deduplication storage

Using a Data Domain system as a target for RMAN backups is relatively straightforward since the system appears as normal disk storage. However, planning for capacity, network throughput, replication bandwidth, recovery operations, and so on is required to ensure that the entire system fulfills the requirements for backup windows, recovery time objectives, desired retention, and support for disaster recovery.

Capacity

Capacity requirements will be determined by the size of the databases being protected, the type of backups being performed, and their retention requirements. A knowledgeable Data Domain system engineer can perform a detailed analysis using tools and models developed over thousands of real-world deployments. Table 2 shows an example of considerations used to size a solution.

Table 2. Sizing consideration example

Database size	450 GB	Combined size of all Oracle datafiles and logfiles
Rate of change	5% per day	Involves 100% change of logfiles and 10% change to database files
Daily backup size	450 GB + 50 GB	Datafiles + archivelogs
Retention period	3 weeks	

While the actual data reduction achieved will be dependent on the amount and nature of the change within the database and logfiles, the capacity chart in Table 3 assumes the combination of deduplication and local (LZ) compression is 40x for each subsequent 500 GB full backup. This means that each daily full backup of 500 GB will result in 12.5 GB of unique data needing to be stored to disk (500 GB / 40 = 12.5 GB).

Table 3. Cumulative storage example

DAY OF WEEK	CUMULATIVE AMOUNT BACKED UP	CUMULATIVE STORAGE CONSUMED	OVERALL COMPRESSION
Sunday	500 GB	125.0 GB	4.0x
Monday	Sunday + 500 – 1000 GB	137.5 GB	7.3x
Tuesday	Monday + 500 – 1500 GB	150.0 GB	10.0x
Wednesday	2000 GB	162.5 GB	12.3x
Thursday	2500 GB	175.0 GB	14.3x
Friday	3000 GB	187.5 GB	16.0x
Saturday	3500 GB	200.0 GB	17.5x
Sunday #2	4000 GB	212.5 GB	18.8x
Sunday #3	7500 GB	300.0 GB	25.0x
Sunday #4	11000 GB	387.5 GB	28.4x

In this example, the user now has three weeks of full backups stored on disk and available for recovery. 11 TB of logical data has been sent to the Data Domain system but due to the effects of deduplication and compression, the actual space consumed is less than 400 GB. At this point, the retention period defined for the backups is reached and RMAN can expire the oldest backup set after the newest backup set is written.

Networking

In theory, the faster the network and the greater the number of network paths, the faster data can move. However, there are other potential bottlenecks in the

communication path that must be considered. For example, due to internal constraints, many clients and backup servers cannot put data onto a network at full line speed. Additionally, each Data Domain system has a rated limit for number of streams and throughput.

As a useful data point, a single gigabit Ethernet connection can move a maximum of 125 MB/s, so to achieve higher aggregate performance on a Data Domain system, more than one link would need to be configured and multiple write streams used. Multiple network connections can also be aggregated to both improve throughput and/or provide additional protection against link failures. Another approach is to use 10 gigabit Ethernet in order to increase the throughput on a single connection.

It is recommended that multiple network connections be configured between the Oracle RMAN server and the Data Domain system to maximize throughput and/or provide redundancy. It is also recommended that these network interfaces be dedicated to RMAN, in order to segregate the backup traffic for administrative or security reasons – this can be as simple as using directly connected network segments between Oracle hosts and the Data Domain system, or as complex as configuring a dedicated VLAN within a multi-tier network infrastructure.

Security

Backup security to a Data Domain system is usually maintained by controlling access to the share points. By limiting hosts to specific directories on the Data Domain system, the DBA can control access to the backup files. By setting UID/GID credentials (in a UNIX environment) or SID credentials (in a Windows Active Directory environment), administrators can add additional controls. In either case, the correct information needs to be provided during setup to ensure that RMAN has access to the necessary path.

It is recommended that the shares (both NFS and CIFS) accessed directly by Oracle RMAN be restricted to the servers actually running Oracle RMAN operations. In Microsoft Windows environments, Active Directory integration will allow consolidated service credentials to be used across the entire deployment, simplifying security management.

Optimizing deduplication

When dealing with any deduplication technology, a few key factors need to be kept in mind. Most important is the fact that any process that modifies a backup data stream on the fly is likely to be detrimental to deduplication. This is because all deduplication techniques rely on being able to identify repeated data patterns. When a backup of otherwise unchanged data is multiplexed, compressed, or encrypted on the fly, the resultant data stream is different each time. This will cause less redundancy to be detected, resulting in the consumption of more physical storage and replication bandwidth for each backup.

Multiplexing, compression, and encryption of backup data were introduced to provide read parallelism as well as enhanced security and performance during backups. Backing up to a Data Domain system is inherently more secure than tape, and does

not suffer performance degradation due to slow backup streams. Thus the recommendation is to disable any of these features when using a Data Domain system as the target for RMAN. Replication traffic can be encrypted to provide security when traversing a WAN, as this is transparent to the replicating Data Domain systems.

Best practices for configuring RMAN options

Selecting the type of backup: Full, incremental, image

It is generally recognized that Oracle recovery processes work best against full backups — a full backup has all of the most recent information and is the fastest to recover. Many Oracle DBAs perform full backups daily as a best practice. When targeted to a Data Domain system, weeks or months of these backups can now be efficiently stored and replicated for disaster recovery. This is the most common deployment scenario at customer sites.

Incremental backups were developed because some databases are just too large to make full backups every night. With deduplication technologies removing redundant copies of database segments, this is now less of an issue. However, a full backup to a Data Domain system still involves first reading all of the data on the server and then moving all of the data over a network connection. Some databases may simply be too large to send over a network in an acceptable period of time. If this is the case, incremental nightly backups are the only available solution with periodic full backups when the opportunity exists, typically on weekends.

To bridge the functional gap between full and incremental backups, Oracle introduced a facility in 10g called Incremental Merge. With this facility, an existing full backup can be updated with data taken during an incremental backup creating, in essence, a new “synthetic” full backup. This benefited the customer by keeping the most recent full copy of the database online on disk. This feature also reduced the RTO time for DBAs to recover the database sooner since the DBA does not need to apply merged backups to the full backup during recovery time. Data Domain enhances this feature by enabling the DBA to keep multiple full copies of their backed up database. A typical scenario would be to be able to keep daily full backups of 30, 60, or 90 days instead of just one full backup.

For databases that cannot be impacted by the RMAN process at all, Oracle presents other data protection techniques that involve storing data on mirrored disk arrays, which can then be separated in order to take backups from the mirror. Oracle also offers standby database solutions where the secondary database can be used to make backups. These advanced solutions are beyond the scope of this paper. However, a Data Domain system still makes an excellent backup repository for even these extremely large databases.

Often overlooked but still critically important is the need to back up and protect the database execution environment as well. An Oracle environment has hundreds of files in addition to the database files themselves that include the executable binaries, shell and Java scripts, and so on. An effective backup scheme should also

include a mechanism for protecting this information. These application environment backups can be located on the same Data Domain system as the database backups to ensure that complete recovery of the Oracle environment can be achieved. Tools such as Oracle Secure Backup, EMC Avamar®, or EMC NetWorker can be effectively utilized to provide this protection.

Logical directory structure for storing backups

If at all possible, it is better to create two backup subdirectories on the Data Domain system for each Oracle instance being backed up — one subdirectory to hold the backup sets comprising the datafiles and the controlfile as well as one to hold the backups of the logfiles. This generally adds one additional step to the RMAN script since “backup database with archivelogs” is no longer a feasible operation and must be separated into separate operations: “backup database” and “backup archivelogs.”

Additionally, it should be noted that logfiles in general do not get very good deduplication. By their nature, they tend to be new and unique each time they are created, therefore it is important to take this into account when evaluating the overall deduplication effect. The preceding recommendation of issuing separate directives to back up the database and the archive logs will allow for separate analysis of the data reduction effect achieved on each. This will enable better capacity management of the Data Domain system over time.

FILESERSET

FILESERSET is the option Oracle RMAN uses to control multiplexing of the datafiles into the backup set, controlling how many datafiles are written to a particular file within the backup set. Since the files involved are read in parallel and asynchronously, the blocks written to the backup set are likely to be in slightly different order each time. While this does not affect the ability to recover the data from the backup set, it does reduce the ability for the Data Domain system to effectively deduplicate the incoming data. Identical data to previous backups may be written in but not recognized as redundant because it is aligned differently or is in a different order.

Specify FILESERSET = 1 or use an image copy when backing up to the Data Domain deduplication storage system.

MAXPIECESIZE

The RMAN parameter MAXPIECESIZE limits the maximum size of a single file in a backup set. In general, it is left at the default setting (which is as large a file as it takes to contain the backup set) but there is one consideration that might suggest restricting it. If Data Domain Replicator software is being used to get the RMAN backup sets moved to the remote DR site, limiting the size of the files in the backup set may allow this replication to be initiated sooner and more frequently. The reason for this is that the replication task on the Data Domain initiates after activity to a particular file has ended. Generally this inactivity time is approximately 10 minutes.

If an administrator wanted to ensure the backup data was getting replicated even more quickly, having several smaller backup set pieces will allow replication to start sooner than having only a few, larger files.

Since replication is often deployed over relatively slow WANs, allowing replication to run concurrently with backups has minimal impact.

By lowering the MAXPIECESIZE, replication will begin sooner. Settings also exist on the Data Domain to throttle the replication throughput, in order to limit the use of a shared WAN. When links are throttled, adjusting this setting may make little to no difference for when replication completes. There is no one value that works for all environments. The choice of settings involves balancing the following concerns:

- Amount of data to be backed up
- Bandwidth of the WAN available for replication
- Any impact on the overall environment due to the larger number of files created with a smaller value specified for MAXPIECESIZE
- Requirement to get the backup data replicated to the DR site as quickly as possible

If the amount of data is small or the need to get data replicated is high, then a smaller MAXPIECESIZE value may be required. If the increased number of files is cumbersome and the WAN link is relatively fast, then a larger setting is suggested. In many cases, replication will complete successfully in the available window when using the default setting.

Parallelism and channels

DBAs can improve throughput by defining additional channels to increase the number of parallel backup processes running. With FILESPERSET=1 (specified per backup job) set to prevent data multiplexing, specifying additional channels and/or degrees of parallelism will often allow RMAN to keep more data moving into the backup sets, minimizing the time required for the backup operation and making best use of otherwise idle system resources.

The optimal number of channels will be determined by the available CPU and memory resources on the Oracle server, as well as I/O capacity of storage where the database being backed up resides. The rate at which a Data Domain system receives data is dependent on the system used, connectivity, and the ability of the host OS to push data over the network. All bottlenecks for a single stream should be identified and resolved where possible prior to adjusting this parameter.

Compression

RMAN provides two styles of compression, unused block compression and binary compression.

The unused block compression is the default for Oracle 10g and can only be disabled by doing an image (COPY) backup. Generally, most customers have been seeing acceptable deduplication when doing “normal” full backups.

To invoke the binary compression of RMAN, the DBA specifies “AS COMPRESSED” on the BACKUP command. This style of compression is similar or better than the hardware compression facility provided by most modern tape backup facilities. However, since this is done in software, it uses scarce CPU and memory resources on the Oracle host during the backup that can be used for other purposes. In many cases, it will significantly slow down the overall performance of an RMAN backup.

Given that the Data Domain deduplication storage system implements a similar compression facility *after* the incoming data has been deduplicated, customers may consider this redundant. As part of an effort to offload the compression function and preserve Oracle CPU resources for driving maximized performance, customers may want to simply allow the Data Domain system to perform the compression.

Encryption

Encryption of backup sets is supported by Oracle RMAN. However, like multiplexing, encryption changes the data patterns presented to the backup device. Customers should understand there is potential performance impact to the deduplication process while leveraging RMAN encryption.

Using multiple archivelog destinations

Oracle can be configured to store archivelogs in multiple locations to provide redundancy and eliminate the need to restore the logs from a backup prior to beginning database recovery operations. A best practice is to use the Data Domain system as a destination for log archiving, as this offers two advantages. First, a second copy of the archivelogs provides additional protection in the event that the primary copy is damaged or lost.

Second, since the archivelogs will eventually be backed up to the Data Domain system as part of an RMAN full backup, this simply means that the files are getting to the Data Domain system sooner. The files are identical (even though RMAN renames them) as long as the FILESPERSET=1 guideline mentioned earlier is followed. The Data Domain system will detect the common data sequences and deduplicate the second copy so that little or no additional disk space is consumed on the Data Domain system. Additional protection without additional disk space consumption is something worth taking advantage of for many environments.

Marking this second destination to Oracle as “OPTIONAL” (as opposed to “MANDATORY”) will allow Oracle to continue running in the event that the Data Domain system is unavailable.

Very large database considerations

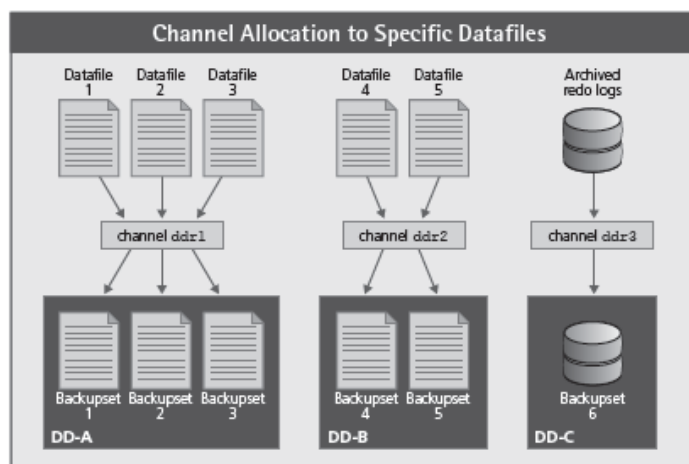
In the case where the database being backed up is consuming tens or hundreds of terabytes, it will be necessary to utilize multiple Data Domain systems to provide

backup resources. This is relatively easy to do, but some planning of the deployment is required.

The Data Domain system can only deduplicate data that it has seen before and to which it has something to compare. Therefore, compression is maximized when the same data is presented to the Data Domain system repeatedly so it has a baseline and can store only the new data. If a large database, composed of possibly thousands of datafiles, is backed up to a series of Data Domain systems, by default RMAN will try to spread the backup load across the various targets. Since this is non-deterministic, RMAN might well send some datafiles to Data Domain system 1 on the first backup but to Data Domain system 2 on a subsequent backup. However, RMAN channels can be used to control the relationship between source datafiles and the targets that the backups are written to. In general, the best practice is to define one or more channels to each Data Domain system and then constrain each channel to backing up specific datafiles, as seen in Figure 6.

This example also shows how the degree of parallelism can be controlled. This RMAN script will run three simultaneous tasks, backing up the specified sets of files (DATAFILE 1,2,3, and 4,5, and ARCHIVELOG ALL) to three different Data Domain systems (specified by the administrator as DD-A, DD-B, and DD-C) simultaneously. This does presume that datafiles 1,2,3, datafiles 4,5, and the archive logs are all stored on different primary disks to minimize disk seeking during the backup.

Another more efficient solution to this is to use a Data Domain Global Deduplication Array (GDA) along with a Data Domain Boost-aware backup application such as NetBackup or NetWorker. These applications can address the GDA and automatically balance the backup load across the two clustered heads deterministically and provide effective backup protection for databases exceeding 200 TB.



```

RUN
{
  ALLOCATE CHANNEL ddr1 DEVICE TYPE DISK FORMAT '/
  ddmnt/DDR-a/backup/oraback';
  ALLOCATE CHANNEL ddr2 DEVICE TYPE DISK FORMAT '/
  ddmnt/DDR-b/backup/oraback';
  ALLOCATE CHANNEL ddr3 DEVICE TYPE DISK FORMAT '/
  ddmnt/DDR-c/backup/oraback';
  BACKUP
  (DATAFILE 1,2,3 FILESPERSET 1 CHANNEL ddr1)
  (DATAFILE 4,5 FILESPERSET 1 CHANNEL ddr2)
  (ARCHIVELOG ALL FILESPERSET 1 CHANNEL ddr3);
}

```

Figure 6. Sample script showing use of RMAN channel parallelism

Conclusion

Getting the most value out of your Oracle data protection strategies requires careful consideration of the benefits and tradeoffs of each approach. Capacity and bandwidth requirements must be balanced with administrative and budgetary constraints to ensure that mission-critical Oracle databases are reliably recoverable within defined service level agreement (SLA) parameters. Data Domain deduplication storage is easily deployed as the critical infrastructure for Oracle data protection best practices, enabling businesses to meet these goals today and in the future.

Deploying EMC Data Domain deduplication storage in Oracle environments delivers the following:

Inline deduplication to enable disk-based data protection for Oracle

- Perform Oracle full backups to disk, reducing the dependence on tape-centric backup infrastructure
- Store weeks or months of full backups using 10-30x reduced disk footprint on average
- Recover Oracle data from full backups, avoiding the complexity and delay of recovering from incremental backups
- Minimize the need for expensive primary disk mirroring and replication technologies for backup data

Seamless integration with Oracle data protection best practices

- Enable direct Oracle backups using simplified RMAN procedures
- Simplify backup monitoring and troubleshooting with common procedures across all Oracle instances

WAN-efficient replication to enable cost-effective disaster recovery

- Avoid the need for tape management as part of an Oracle disaster recovery (DR) deployment
- Improve RPO and RTO service levels by moving data to offsite locations quickly, and automatically, using up to 99 percent less bandwidth
- Simplify and automate Oracle DR procedures utilizing the replicated copies of all RMAN backups
- By following the best practice recommendations presented in this paper, customers can maximize the benefits of using Data Domain systems to protect Oracle databases using RMAN, and ensure a smooth integration into their environment.

References

For additional information, see the following:

- EMC Data Domain Family products and deduplication technology
<http://www.emc.com/products/family/data-domain-family.htm>
- EMC Data Domain solutions for Oracle
<http://www.datadomain.com/solutions/oracle.html>
- EMC Backup, Recovery, Archive solutions for Oracle
<http://www.emc.com/solutions/samples/oracle/backup-recovery-archive-oracle.htm>
- EMC Solutions for Oracle
<http://www.emc.com/solutions/application-environment/oracle/index.htm>
- EMC Data Domain Global Deduplication Array
<http://www.emc.com/products/detail/hardware/data-domain-global-deduplication-array.htm>
- EMC Data Domain Boost software
<http://www.emc.com/products/detail/software/data-domain-boost.htm>
- *EMC Data Domain SISL Scalability Architecture – A Detailed Review* white paper
<http://www.emc.com/collateral/hardware/white-papers/h7221-data-domain-sisl-sclg-arch-wp.pdf>
- *EMC Data Domain Replicator – A Detailed Review* white paper
<http://www.emc.com/collateral/software/white-papers/h7082-data-domain-replicator-wp.pdf>
- *EMC Data Invulnerability Architecture: Ensuring Data Integrity and Storage System Recoverability* white paper
<http://www.emc.com/collateral/software/white-papers/h7219-data-domain-data-invul-arch-wp.pdf>