

# **DATABASE BACKUP & RECOVERY**

## **STRATEGIES AND BEST PRACTICES**

*Tammy Bednar, Oracle Corporation*

### **INTRODUCTION**

A 500-gigabyte database used to be known as a large database. Today multi-terabyte databases are the norm – and the trend will continue in the foreseeable future. With a Terabyte of database storage, one could store 500 movies online, to watch at the click of a button. A concern with such large database is the time it takes to backup all of the data. I hear quite often, “I need a fast backup”. My reply is, “What is the requirement to restore?” We seem to focus so much on the ability to backup quickly that we forget that the reason to backup is to ensure the capability to recover your data.

Recovering database data from a backup involves three distinct operations: restore the data, roll the backup forward to a more recent time by applying redo data, and roll back all changes made in uncommitted transactions to their original state. In general, recovery refers to the various operations involved in restoring, rolling forward, and rolling back a backup. Backup and recovery refers to the various strategies and operations involved in protecting the database against data loss and reconstructing the database should a loss occur.

This paper highlights Oracle database backup and recovery best practices. First, and most importantly, is to create a backup and recovery plan. The key to a successful recovery begins with a backup. Backups can be made with Oracle’s Recovery Manager (RMAN) or third-party tools. Second, I will examine how to decrease the recovery time by implementing strategies to keep important data readily available, making tablespaces read-only, and the importance of backing up the control file. Finally, I have highlighted RMAN best practices to backup up Real Application Clusters, RMAN performance, and logical naming of backups.

### **THE KEY TO SUCCESSFUL RECOVERY**

#### **BACKUP AND RECOVERY PLAN**

Backup and recovery is one of the most important aspects of database administration. Whether companies operate a single database or multiple databases storing hundreds of gigabytes or even terabytes of data, they share one common factor — the need to back up important data and protect themselves from disaster by developing a recovery strategy. Time, resources and effort involved in designing a recovery plan deter many from taking steps to protect their businesses from disaster. In fact, only 45 percent of Fortune 500 enterprises have formal business recovery plans in place. Of those, only 12 percent are considered effective at an enterprise level.<sup>1</sup>

A database contains a wide variety of types of data. When developing a backup strategy, DBAs must decide what information they want to protect. In deciding what to back up, the basic principle is to prioritize data depending on its importance and the degree to which it changes. Planning for a data loss will prepare your enterprise for the unexpected outages that can occur in the day-to-day operations or disaster recovery. These unexpected outages may fall under categories such as hardware failures, software failures, human errors and acts of nature. Documenting the

---

<sup>1</sup> “Calculating the Cost of Downtime”, Angela Karr, Advanstar Communications

current environment will help you and your administrators assess the recoverability of your data. If a failure or disaster does occur, the database administrator is able to diagnose and quickly correct the situation using this recovery planning documentation.

Figure 1 below represents a template to take inventory of the hardware components and database configurations. Detailed database configurations should also include the backup and recovery methods.

Hardware Configuration	Database Configuration
Vendor/Model	Instance Name
Operation System	Host Name
Version/Patch release	RDBMS Version
Disk Capacity	Size of Database
No. of Disk/Controller	Backup Method/Frequency
Availability Requirement	Backup Method/Time to Restore
Media Mgmt Vendor	Datafile mount point(s)
Type and no of tapes	

**Figure 1 Template to document for your backup and recovery plan.**

### *EFFECTIVE BACKUP AND RECOVERY VS. THE COST*

There are many strategies and architecture solutions to protect your data and the ability to recover in a timely manner is based on the needs of the business and the cost to implement the solution. The cost of recovery ranges from a low-end solution with a large recovery time to the high-end solution with a rapid recovery time. It may include the simple option of backing up files to tape for offsite storage with a relatively large recovery time or host storage-based snapshots for more rapid recovery but no site failure protection. Here is a list of common solutions to protect your data.

- RAID technology
- Take frequent backups
- Keep “*redundancy set*” online for quick access and fast recovery
- Take snapshots using host or storage-based snapshot tools
- Use Oracle multiplexing to create copies of critical files
- Integrated solution with RMAN
- Protect with a standby database

When you choose your recovery strategy and your choice is one of the above, either by itself or in combination with another technique, your odds of losing your data are greatly reduced. But, a backup is still the key to a successful recovery.

### IS DOWNTIME ACCEPTABLE?

Oracle database backups can be made while the database is open or closed. Planned downtime of the database can be disruptive to operations, especially in global enterprises that support users in multiple time zones, up to 24-hours per day. In these cases it is important to design a backup plan to minimize database interruptions.

Merrill Lynch, one of the world's leading financial management and advisory companies<sup>2</sup>, requires 24x7 uptime Monday thru Friday. On Saturday and Sunday, only 6 hours of uptime is required. The other 18 hours of the day can be used for planned outages to upgrade applications and perform hardware maintenance.

Depending on your business, some enterprises can afford downtime. If your overall business strategy requires Five Nines, then your backup strategy should implement an online backup. The database needs never to be taken down for a backup.

How Many 9's		Maximum Downtime per Year
One Nine	90.000%	36 days
Two Nines	99.000%	3.7 days
Three Nines	99.900%	9 hours
Four Nines	99.990%	53 minutes
Five Nines	99.999%	5 minutes

**Figure 2 How many nines do you require?**

### WHAT IS YOUR MTTR?

What is your mean time to recover (MTTR)? There are two types of recovery that you may have to perform on an Oracle database. The first type is called instance recovery. Instance recovery occurs when the database processes stop unexpectedly. The database cannot operate if a hardware disk fails or a controller ceases to function. If the database crashes, do you know how long does it takes for the database to open and become productive again?

Oracle9i offers fast-start fault recovery functionality to control instance recovery. This reduces the time required for cache recovery and makes the recovery bounded and predictable. Many service level agreements (SLA) require the IT staff to reliably set a target on the time it will take to recover the database.

Administrators specify a bounded time to complete the cache recovery phase of recovery with the FAST\_START\_MTTR\_TARGET initialization parameter. The FAST\_START\_MTTR\_TARGET initialization parameter lets you specify in seconds a target mean time to recover. The database server automatically adjusts the write rate to meet the specified recovery target.

The second type of recovery a DBA may perform is media recovery. Media recovery is required anytime a restoration of any database file is required. The factors that affect the MTTR of media recovery include the backup media location, the time to restore the file from the backup media, and the re-application of transactions for that file. Keeping these factors in mind, does your restore plan meet the SLA?

### CAN YOU AFFORD TO LOSE ANY DATA?

Database logging may ensure that no transactions are lost if the database must be recovered from a previous backup. To decide if logging is required for a database, ask yourself, "Can I afford to lose any data?" Every time a change in the database occurs, Oracle generates a record of the transaction into the online redo log, which is on disk. Choosing

<sup>2</sup> Merrill Lynch is one of the world's leading financial management and advisory companies with offices in 44 countries and total client assets of about \$1.6 trillion. As an investment bank, Merrill Lynch is the top global underwriter and market maker of debt and equity securities and a leading strategic advisor to corporations, governments, institutions, and individuals worldwide. Through Merrill Lynch Investment Managers, they're one of the world's largest managers of financial assets.

to keep these change before they are overwritten is called archiving. If you cannot afford to loose data, your backup plan must include the ability to backup archive logs.

Archived redo logs are crucial for recovery when no data can be lost, since they constitute a record of changes to the database. Oracle can be run in either of two modes:

- ARCHIVELOG -- Oracle archives the filled online redo log files before reusing them in the cycle.
- NOARCHIVELOG -- Oracle does not archive the filled online redo log files before reusing them in the cycle.

Running the database in ARCHIVELOG mode has the following benefits:

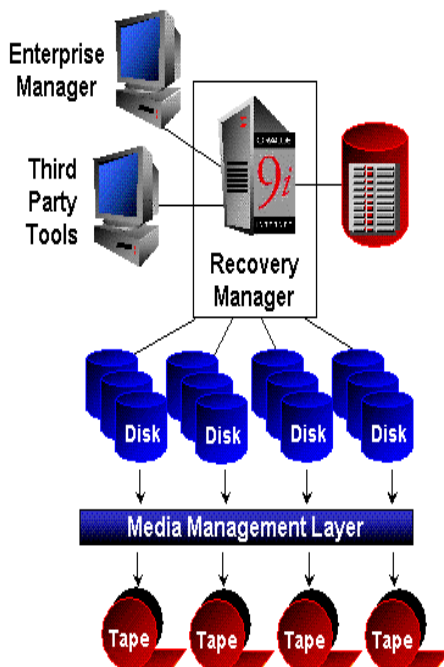
- The database can be completely recovered from both instance and media failure.
- The user can perform backups while the database is open and available for use.
- Archived redo logs can be transmitted and applied to the physical standby database, which is an exact replica of the primary database.
- Oracle supports multiplexed archive logs to avoid any possible single point of failure on the archive logs.
- The user has more recovery options, such as the ability to perform tablespace-point-in-time recovery (TSPITR)

Running the database in NOARCHIVELOG mode has the following consequences:

- The user can only back up the database while it is completely closed after a clean shutdown.
- Typically, the only media recovery option is to restore the whole database, which causes the loss of all transactions since the last backup.

## BACKUP TOOLS AND METHODS

To backup the Oracle database, there are two types of tools that can be used. Oracle offers Recovery Manager free of charge. Recovery Manager is Oracle's utility to manage the backup, and more importantly the recovery, of the database. It eliminates operational complexity while providing superior performance and availability of the database. The second tool to backup the database is any third-party application in which the user must manage the backups of the database or the traditional method of writing custom, homegrown scripts which may require constant update.



### RECOVERY MANAGER

Recovery Manager (RMAN) debuted with Oracle8 to provide DBAs an integrated backup and recovery solution. Recovery Manager determines the most efficient method of executing the requested backup, restore, or recovery operation and then executes these operations in concert with the Oracle database server. Recovery Manager and the server automatically identify modifications to the structure of the database and dynamically adjust the required operation to adapt to the changes.

Oracle9i Recovery Manager features enhance database availability and manageability by providing the capability to recover corrupted database blocks while the datafile remains online, create persistent backup/restore configurations which provides consistent operation's behavior, and automatic control file backup and restore to protect against disaster.

There are a number of significant benefits to using Recovery Manager.

- Automated database backup and recovery

Figure 3 Recovery Manager Architecture

- Allows backing up of the entire database or any logical unit such as the controlfile, datafile, tablespace or archive log files
- Incremental backups
- Backups can be performed while the database is open or closed
- Two types of backup: image copies or backup sets
- Intelligent archive log management for both backup and recovery
- Tablespace Point-In-Time Recovery support
- Integration with the Oracle Enterprise Manager Backup Manager GUI
- Omission of empty blocks during backup for optimization
- No extra redo is generated during online backups reducing any potential for a performance penalty

### *ENTERPRISE MANAGER SUPPORT*

A Recovery Manager Graphical User Interface is built into Oracle's Enterprise Manager. Enterprise Manager is the central management framework providing a robust console, a rich set of tools, and the extensibility to detect, solve, and simplify the problems of any managed environment. Enterprise Manager is included as part of the Oracle Database Server.

Scheduling of backup jobs and RMAN-specific tasks via the Enterprise Manager Job System enhances the Backup and Recovery facility and provides the flexibility to schedule the backup tasks at specified intervals, on specified day of the week, or on a specified day of the month.

### *MEDIA MANAGEMENT INTEGRATION*

Organizations rely on Oracle to provide solutions for very large critical systems. In addition to needing databases capable of handling large amounts of data and complex queries, these organizations also need robust backup and recovery technology. Recovery of data quickly and reliably is paramount should some aspect of the system fail. To address these needs, Oracle has created the Backup Solutions Program (BSP), a cooperative program designed to facilitate tighter integration between Oracle's backup products and those of third-party media management vendors.

Why learn two tools to backup your enterprise systems? Use your preferred third-party backup software to backup all of your system files and the Oracle database. Use a product that is integrated with Recovery Manager to ensure your database is precisely backed up; "One tool – One interface". Together, Oracle and media management vendors provide robust easy-to-use database backup and recovery solutions to customers with high-end requirements. To view the current members of the Backup Solutions Program, go to <http://otn.oracle.com/deploy/availability>.

### *USER-MANAGED BACKUPS*

Any third-party tool or manually created script used to backup the database without RMAN is called user-managed backup. It is the responsibility of the DBA to ensure backups have been made of all database files. If the database is logging all changes, the archive logs also require to be backed up. In addition, the DBA or third-party tool must implement some type of catalog to track the backup history. If restoration of the database is required, then the DBA must select which backup to restore as well as any archive logs needed to roll forward database transactions.

## BACKUP METHODS

Backups are divided into physical and logical backups. Physical backups are copies of the database files. The phrase “backup and recovery” refers to the transfer of copied files from one location to another, along with the various operations performed on these files. Restoring a physical backup means reconstructing it and making it available to the Oracle server. To recover a restored backup, data is updated using the change records from the online log. The online log records changes made to the database after the backup was taken.

In contrast, logical backups contain data that is exported using SQL commands or Oracle Export utility. Logical backups can be used to supplement physical backups, but the primary use of a logical backup is to move data between heterogeneous hosts.

Recovery Manager offers two types of backup: RMAN backup sets and image copies. The Recovery Manager backups, which are proprietary structures, can only be restored using RMAN. The advantage of RMAN type backups is they only contain blocks that have ever been used in the database. If the database contains unused space, the storage of the backups is smaller than the total size of the database.

Image copy backup is a mirror copy of the file, as it exists during backup. The advantage of making an image copy backup is that it can be used immediately without any tool required to uncompress it. Image copies of the Oracle database can also be made by OS utilities such as dd or copy.

A combination of methods and tools will ensure that the Oracle database is recoverable. A Merrill Lynch DBA says “...the key [to recovery] is to ensure you have the archived logs backed up at least twice and on two different backup tapes before you delete them. It would be a real shame to have a tape wreck and then be unable to recover the database. You can never have too many backups of the archived redo log files.” Table 1 compares backup and recovery features using Recovery Manager and OS utilities.

Backup and Recovery Feature	Recovery Manager	User Managed
Closed database backups	Supported.	Supported.
Online database backups	Server Managed.	Use BEGIN/END BACKUP statements.
Incremental backups	Supported.	Not Supported.
Corrupt block detection	Supported..	Not Supported.
Automatic backup	Supported..	Not Supported.
Backup catalogs	Supported. Backups are automatically recorded in the control file and optionally in the recovery catalog.	Not supported. The user must manually maintain a backup catalog.
Backups to media manager	Supported. Interfaces with a media manager. RMAN also supports proxy copy, a feature that allows the media manager to manage the transfer of data	Supported. Backup to tape is manual or controlled by a media manager.
Backs up parameter file and password files	Supported.	Supported.

**Table 1 Comparison of Recovery Manager and User Managed Methods**

## TESTING THE RECOVERY PLAN

Several problems can halt the normal operation of an Oracle database or affect database IO operations. You can never be too ready to recover from a failure. Performing regular test recoveries ensures that your backup is working and it also helps you stay familiar with recovery procedures. If recovery is required, additional resources may be needed and recovery time extended. An important part of the backup and recovery plan is to document the processes required to complete a successful restoration so, if necessary, the non-primary DBA can execute it. Schedule regular testing to perform recovery-using examples of known failures and disasters.

Some of the most common type of database problems are media failures, block corruptions, user errors, and disasters. For some of these problems, crash and instance recovery occur automatically and require no action on the part of the database administrator. For other problems, administrator-initiated media recovery is required.

### *MEDIA FAILURE*

An error can occur when trying to write or read a file on disk that is required to operate an Oracle database. This is called a media failure because there is a physical problem reading or writing to files on the storage medium. A common example of media failure is a disk head crash that causes the loss of all database files on a disk drive. All files associated with a database are vulnerable to a disk crash, including datafiles, control files, online redo logs, and archived logs.

The appropriate recovery from a media failure depends on the files affected. Media failure is the primary concern of a backup and recovery strategy, because it typically requires restoring some or all database files and the application of redo during recovery.

### *BLOCK CORRUPTION*

What types of corruption can occur within an Oracle database block? Faulty hardware or an operating system bug can cause an Oracle block that is not in a recognized Oracle format, or whose contents are not internally consistent. Oracle identifies corrupt blocks as one of two types:

- Logically corrupt. For example, the block was corrupted by an incorrect block type but does not appear to be media corrupt.
- Media corrupt, that is, the block format is not correct. The block may have:
  - An incorrect checksum
  - A wrong data block address
  - An invalid block type

Oracle's Recovery Manager feature, Block Media Recovery, is a technique for restoring and recovering an individual corrupt datablock or set of datablocks within a datafile. Although datafile media recovery is the principal form of recovery, you can use the Oracle9i Recovery Manager interface to perform block media recovery for cases when a small number of blocks require media recovery. Block media recovery provides the several advantages over datafile media recovery.

- Lowers the Mean Time to Recovery (MTTR) because only blocks needing recovery are restored and only necessary corrupt blocks undergo recovery.
- Block media recovery minimizes redo application time and avoids IO overhead during recovery.
- Allows affected datafiles to remain online during recovery of the blocks. Without block-level recovery, if even a single block is corrupt you must restore a backup of the entire datafile and apply all redo generated for that file after the backup was created.

### *USER ERROR*

As an administrator, you can do little to prevent user errors such as accidentally dropping a table. Often, increased training on database and application principles can reduce user error. You can also avoid user errors by administering privileges correctly so that users are able to do less potential damage. Furthermore, by planning an effective recovery scheme ahead of time, you can ease the work necessary to recover from user errors.

Typically, a user error such as a dropped table requires re-entering the lost changes manually, if a record of them exists, importing the dropped object, if an export file exists, or performing incomplete recovery either of an individual tablespaces or of the entire database. If a standby database is in place, the table may be exported from the standby database or failover to the standby before the change has been applied.

### *DISASTER*

When your hardware host or worse, the entire facility is lost, your recovery plan must include provisions to recover at an off-site host. One of the first decisions to be made is getting your backups away from the primary site. Merrill Lynch has set up a backup facility 16 kilometers from their main facility. The backups of the database and system occur over a gigabit network so that manual shipping of backup media for vaulting is not required.

Once the backups are located offsite, the next challenge is to restore the system and database to the host. A DBA that works for a major bank once told me that her bonus is tied to the ability to recover the database to meet their SLA. In order to achieve that, their off-site testing occurs very often. When a fire drill is scheduled, most people can be very calm. To truly test out a disaster recovery, an unscheduled fire drill can indicate if your recovery plan is sound.

## **BACKUP DISK CACHE**

It is a fact that recovering from disk is always faster than recovering from tape. If the Oracle database is in archivelog mode, most customers have implemented a strategy to keep a certain number of archive logs to enable recovery to a point in time of their choosing. Merrill Lynch keeps the archive logs on disk for seven days. In that seven-day time frame, the archive logs are backed up to tape twice exclusively using Legato<sup>3</sup> Systems' NetWorker product and Oracle module, which interface with RMAN. In addition to the RMAN backups, Merrill Lynch uses Legato NetWorker® to backup the directory in which the archive logs reside. After the second backup, the archive logs are then compressed to make room for newly created archive logs. After seven days, the archive logs are deleted from disk.

When more disk space is devoted for recovery, more than archive logs can be put into the backup disk cache area. Backups of the database can be made directly to disk using RMAN, OS utilities, and media-manager tools. You may want to consider making image type backups of datafiles since it is fast operation to use the backup image copy when a file may require recovery.

## **TO BACKUP OR RELOAD - THAT IS THE QUESTION**

If you cannot afford to lose any transactions, then the database should be in archivelog mode. In a data warehousing environment or in the case where a batch data loading occurs, archive logs can fill up a disk quickly. Oracle offers a table level feature called NOLOGGING, which minimizes the information written to redo logs. Processes running with the NOLOGGING option set run faster because no redo is generated. The redo records contain enough information to determine that an object has data inserted into it, but not enough to recover the object using the archive logs should media recovery be necessary. After a NOLOGGING operation against a table, partition, or index, if a media failure occurs before a backup is taken, then all tables, partitions, and indexes that have been modified are not recoverable.

---

<sup>3</sup> Legato System, Inc. (NASDAQ: LGTO) directly and through strategic partnerships and alliances, develops and delivers the software solutions and services that provide for business continuance by leveraging the advantages of enterprise automation.

Amazon.com has a database that loads many gigabytes of data hourly. Their backup and recovery strategy uses RMAN to make an online backup of the database every Sunday since there is very little or no data loading occurring. Since Amazon.com is also using the NOLOGGING-direct path loading option, they know that media recovery using just the archive logs may render some objects unrecoverable. The database is so large and there is not enough tape drive bandwidth that a backup during the week cannot be completed during the time when no data is being loaded. Amazon.com keeps a week's worth of data that was loaded. If recovery is necessary, Amazon.com restores the previous Sunday's backup and opens the database. Then the data is reloaded from the week's worth of data files.

## **OFFLOAD THE BACKUPS**

Oracle has been providing customers with data protection solutions in the form of standby database technology since release 7 of the database server product. Over time, the scope of what was once thought of only as a disaster recovery solution has broadened, to protect data from all threats. Starting with the release of Oracle 8.1.7, RMAN can be used to backup the standby database to offload any overhead from the production environment. The application of archive logs to the standby continues without interruption.

RMAN backs up the standby database and its associated archived redo logs. Standby backups of datafiles and archived redo logs are fully interchangeable with primary database backups. In other words, you can restore a backup of a standby datafile to the primary database, and you can restore a backup of a primary datafile to the standby database. Backing up standby files is often better than backing up the production files, for the following reasons:

- Offload processing from production host.
- Because the standby database is not the production database and a standby backup does not interfere with transactions or batch jobs in the production database.
- If the standby and primary databases are on separate hosts, then standby backup operations do not consume CPU cycles, allocate memory, or consume other resources on the production host.

## **WHY DO I NEED TO BACKUP THE CONTROL FILE?**

Every Oracle database has a control file. A control file is a small binary file that records the physical structure of the database. The control file must be available for writing by the Oracle database server whenever the database is open. Without the control file, the database cannot be mounted and recovery is difficult. This important file also contains consistency information that is used during recovery, such as the:

- Database name
- Timestamp of database creation
- Names of the database's datafiles and online and archived redo log files
- Checkpoint, a record indicating the point in the redo log where all database changes prior to this point have been saved in the datafiles
- Recovery Manager (RMAN) backup meta-data

Every time a user mounts an Oracle database, its control file is used to identify the datafiles and online redo log files that must be opened for database operation. If the physical makeup of the database changes, a new datafile or redo log file is created, Oracle modifies the database's control file to reflect the change. Users can multiplex the control file, allowing Oracle to write multiple copies of the control file to protect it against disaster. If the operating system supports disk mirroring, the control file can also be mirrored, allowing the O/S to write a copy of the control file to multiple disks.

The control file should be backed up whenever the structure of the database changes. Structural changes can include adding, dropping, or altering datafiles or tablespaces and adding or dropping online redo logs. Recovery operations are less error prone when using the database's current control file.

If Recovery Manager is integrated into your backup and recovery strategy, the control file is automatically backed up anytime a backup or recovery operation is executed. This eliminates the need to manually backup up and track versions of the file. When necessary, RMAN automatically restores the valid control file version during a recovery operation.

### **BACKING UP READ ONLY TABLESPACES**

Recovery is not needed for read-only tablespaces during crash or instance recovery. Making a tablespace read-only prevents write operations on the datafiles in the tablespace. The primary purpose of read-only tablespaces is to eliminate the need to perform backup and recovery of large, static portions of a database, but they also provide a means of completely protecting historical data so that no one can modify the data after the fact. Since read-only tablespaces can never be updated, they can reside on CD-ROM or WORM (Write Once-Read Many) devices.

Because read-only tablespaces cannot be modified, they do not need repeated backups. Recovery may not be necessary if the tablespace is intact. In a large table that has partitioned the data, you may want to consider creating a partition per tablespace. Once the data in the partition has become static, make the tablespace read-only to decrease the number of times it requires to be backed up.

### **RECOVERY MANAGER BEST PRACTICES**

Recovery Manager (RMAN) is Oracle's utility to manage the backup, and more importantly the recovery, of the database. To use RMAN in the most efficient manner, let's take a look at backing up Real Application Clusters, performance tuning, the retention policy, tagging backups, and the recovery catalog

### **BACKING UP REAL APPLICATION CLUSTERS**

Real Application Clusters harnesses the processing power of multiple, interconnected computers. Real Application Clusters software and a collection of hardware, known as a cluster, unite the processing power of each component to become a robust computing environment. In Real Application Clusters environments, all active instances can concurrently execute transactions against a shared database. Real Application Clusters coordinates each instance's access to the shared data to provide consistency and integrity.

Backing up a Real Application Clusters environment does not differ substantially from a single-instance environment. The consideration that must be made is setting up the archive log directories and how the other nodes of a cluster can access them. If a clustered file system is used for a Real Application Clusters database, any node can read/write to the directories of the archive logs. If a clustered file system is not available, then below are three schemes to backup Real Application Clusters archive logs with Recovery Manager. For additional archive schemes, refer to the Real Application Clusters documentation.

### SHARED READ LOCAL ARCHIVING

In the shared read local archiving scheme, each node writes to a single local archived log destination and can read the archived log files of the other nodes. Read access is commonly achieved using NFS on UNIX or shared drives on Windows platforms.


#### ADVANTAGES AND DISADVANTAGES OF THE SHARED READ LOCAL ARCHIVING SCHEME

##### Advantages

- None of the nodes archives logs over the network.
- You can back up all logs from any node in the cluster.
- If each node has a local tape drive, then you can distribute an archived log backup so that each node backs up local logs without accessing the network.
- Recovery is simplified because every node has access to all archived logs.

##### Disadvantages

- The setup is complex because each node must have NFS read access to the other nodes in the cluster.
- This scheme has a single point of failure. If one node fails after the most recent complete backup, then the archived logs located on this node are unavailable to the surviving nodes for recovery.
- If only one node has a local tape drive, then an archived log backup must read logs through the NFS connections.



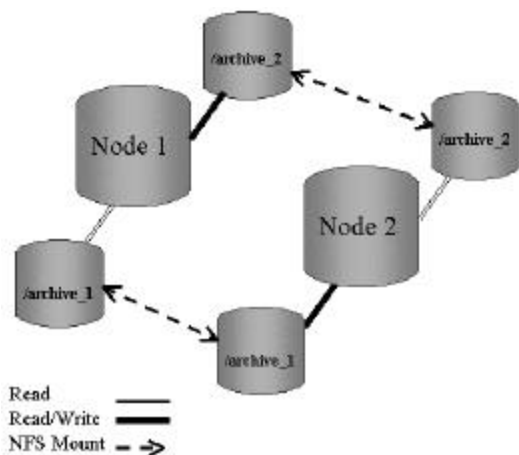
### SHARED READ, ONE REMOTE ARCHIVING

A shared read, one remote archiving scheme in which each node writes to both a local destination and to a remote destination on one other node. Also, each node has read-only access to a directory on one of the other nodes.

#### ADVANTAGES AND DISADVANTAGES OF THE SHARED READ, ONE REMOTE SCHEME

##### Advantages

- You can perform media recovery from any node in the cluster.
- You can back up all logs from any node in the cluster without performing manual transfers.
- If each node has a local tape drive, then you can distribute an archived log backup so that each node backs up local logs without accessing the network.
- Only one node needs to be available in order to perform media recovery.
- An upgrade to a clustering file system does not require changes in the Oracle configuration.



##### Disadvantages

- The setup is complex because each node must have NFS access to the other nodes in the cluster.
- Each node archives to the other nodes in the cluster through a network connection.
- If secondary locations are set as MANDATAORY, then production can stop if the network goes down.

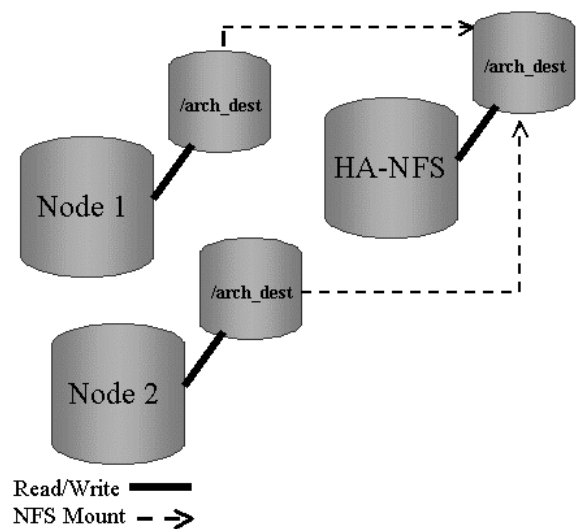
### ARCHIVING TO A CENTRAL NFS DIRECTORY FOR HIGH AVAILABILITY

Configuring a central high availability NFS (HA-NFS) directory requires that each node mount the same destination device for archiving.

#### ADVANTAGES AND DISADVANTAGES OF THE CENTRAL NFS ARCHIVING SCHEME

##### Advantages

- You can perform media recovery from any node in the cluster.
- You can back up all archived logs from any node in the cluster without performing a manual transfer.
- If only one node has a local tape drive, then you can specify `BACKUPARCHIVELOG ALL DELETE INPUT` to delete all logs from from the HA-NFS server after backing them up.
- Only one node needs to be available in order to perform media recovery.
- An upgrade to a clustering file system does not require changes to the Oracle configuration.
- Because every node archives to a single directory that is named the same on every node, the complexity of the configuration does not increase significantly if you more nodes to the cluster.



##### Disadvantages

- You must maintain an additional node.
- Because Oracle must be able to archive every thread to at least one directory, if the network goes down, then the database stalls the next time it attempts to archive a log.
- Each node must archive its logs over the network.

### BACKUP AND RESTORE PERFORMANCE

Whether companies operate a single database or multiple databases that store hundreds of gigabytes or even terabytes of data, they share a common factor: the need for a fast and reliable backup solution. When evaluating RMAN as a backup and recovery solution, many users ask the following questions:

- What is the optimal RMAN backup setting?
- How is the performance of an RMAN backup affected by the alteration of backup parameters?

Below is a summary of performance factors that affect RMAN backups and recoveries. To view an in-depth report on RMAN performance, please go to <http://otn.oracle.com/deploy/availability>.

## PERFORMANCE FACTORS

Tuning an RMAN backup does not have to be complex and labor intensive. You can achieve the backup and recovery performance for your enterprise systems by examining your disk storage configuration, use of asynchronous IO, allocation of RMAN channels, file multiplexing, and disk and tape buffers.

- *DISK STORAGE CONFIGURATION*

Configuring storage subsystems for the Oracle database is an unnecessarily complex process. To achieve high availability the disks should be mirrored. A simple, efficient, and highly available storage configuration is possible by making extensive use of striping across large sets of disks. We call this methodology **S.A.M.E.**<sup>4</sup> SAME stands for Stripe and Mirror Everything. Using SAME, with one-megabyte stripes, RMAN backup test results show that reading striped disks is significantly more efficient than reading nonstriped disks. The backup is at least 10% faster while the CPU utilization is much smaller.

- *ASYNCHRONOUS IO*

When Oracle reads or writes data to disk, the IO can be either synchronous or asynchronous. When IO is synchronous, a server process can perform only one task at a time. On the other hand, when IO is asynchronous, a server process can begin an IO and then perform other work while waiting for the IO to complete. The ability to do asynchronous IO improves performance.

Native asynchronous IO can be enabled with the initialization parameter `DISK_ASYNC_IO`. On operating systems that do not support native asynchronous IO, Oracle can simulate it by using disk IO slave processes that are dedicated to performing IO on behalf of another process. IO slave processes used by the ARCH, LGWR, and backup processes are configured with the `DBWR_IO_SLAVES` initialization parameter.

- *RMAN CHANNEL PARALLELISM*

An RMAN channel represents one stream of data to an output device. When RMAN allocates a channel, it establishes a connection to a target database instance by starting a server session on the instance. This Oracle server session performs the backup, restore, and recovery operations. So, the RMAN channel is an operating system process or thread that reads the data from disk and sends it to the output. The backup is done in parallel on all allocated channels. Because the channels act independently, the number of active channels defines the level of parallelism. When backing up the database, only allocate one channel for each output tape device or disk subsystem.

- *BACKUP MULTIPLEXING*

A RMAN backup set, which is a logical object, contains one or more physical backup pieces. By default, one backup set contains one backup piece. Backup pieces contain the backed up datafiles, control files, or archived redo logs. A backup set is one or more backup pieces that make up a full or incremental backup of the objects specified in the backup operation.

The technique of RMAN multiplexing is to simultaneously read files on disks and then write them into the same backup piece. For example, RMAN can read from 2 datafiles simultaneously, and then combine the blocks from these datafiles into a single backup piece. The level of multiplexing is the number of files read simultaneously on a single channel and then written to backup buffers.

Multiplexing may add more time to a restore operation since more than one file is intertwined in one RMAN backup set. To eliminate this additional restore time, set `FILESERSET` to the value between 1 or 2 when using SAME, otherwise 4 to 8.

- *ASYNCHRONOUS AND SYNCHRONOUS IO FOR DISK AND TAPE*

---

<sup>4</sup> For more information on S.A.M.E, go to <http://otn.oracle.com/deploy/availability>.

Even though tapes are sequential devices, writing to and reading from a tape can be asynchronous. Oracle simulates asynchronous tape IO by using backup tape IO slave processes. These slave processes are enabled with the `BACKUP_TAPE_IO_SLAVES` initialization parameter.

Oracle backups use two different types of buffers for reading and writing of data: disk buffers and tape buffers. Disk buffers are used for reading and writing to disks, whereas tape buffers are used for tape IO operations. The size of these buffers defines the amount of data transferred in a single IO call. In asynchronous IO, operations are initiated without waiting for the first operation to complete. In Oracle9i, the default tape buffer size is 256 KB. The sizes of disk buffers are determined dynamically.

- *MEDIA MANAGEMENT SOFTWARE*

Some media management products can compress data. Real-world experience shows that software compression does not help improve local backup performance. It helps only if the data is transferred over a slow network. In some cases the media management software may be misconfigured so that the transfer of data to a backup media is very slow. For example, the data is copied by means of network protocols such as TCP/IP even though the tape is locally attached.

The physical tape block size can affect backup performance. The block size is the amount of data written by media management software to a tape in one write operation. The common rule is that a larger tape block size leads to a faster backup. For more efficient backup to tape, set the RMAN parameter `BLKSIZE` to be greater than or equal to 256 kilobytes. The tape block size and the RMAN `BLKSIZE` parameter should be equal.

- *DB\_FILE\_DIRECT\_IO\_COUNT*

The sizes of the disk and tape buffers should be relatively large. For Oracle8i, set the disk buffer size to 1 megabyte by changing the initialization parameter `DB_FILE_DIRECT_IO_COUNT`. For Oracle9i, this change is not required since it has been updated for you.

## **BACKUP RETENTION POLICY**

Recovery Manager implements a retention policy that governs when backups expire. Recovery Manager automatically marks as obsolete all backups and archived logs no longer required to recover the database. Implementing a backup retention policy features is designed to reduce the time and effort administrators spend in performing routine administration through automation of the most commonly performed tasks.

The retention policy is configured using a recovery window or by specifying backup redundancy. The recovery window is a period of time in days extending backwards from the present. Recovery Manager will ensure database will be recoverable to any point of time within that number of days. Redundancy specifies the number of backups, or copies, should be retained to ensure recoverability. When integrating with a third-party media management vendor, the retention policy should be governed by RMAN and not by the tape recycle policy. Since backups can span several tapes, you may risk lose half of backup if the tapes are not recycled at the same time, or worse, you may lose the capability to restore from a past backup.

## **TAG, YOU'RE IT**

A tag is a symbolic name for a backup set or file copy such as `WEEKLY_BACKUP`. You can assign a user-specified character string called a tag to backup sets and image copies. You can specify the tag rather than the filename when restoring a specific backup. Tags do not need to be unique, so multiple backup sets or image copies can have the same tag. When a tag is not unique, the tag refers to the most current suitable file. By default, RMAN selects the most recent backups to restore unless qualified by a tag. The most current suitable backup containing the specified file may not be the most recent backup, as can occur in point-in-time recovery. Tags can indicate the intended purpose or usage of different classes of backups or file copies.

## **RMAN RECOVERY CATALOG**

The RMAN repository is a collection of metadata about the target databases that RMAN uses to conduct its backup, recovery, and maintenance operations. You can either create a recovery catalog in which to store this information, or let RMAN store it exclusively in the target database control file. Although RMAN can conduct all major backup and recovery operations using just the control file, some RMAN commands function only when you use a recovery catalog. For example, RMAN scripts can only be stored in a recovery catalog or view the database structure, as it existed in the past.

Never store a recovery catalog containing the RMAN repository for a database in the same database or on the same disks as the target database. For example, do not store the catalog for database prod1 in prod1. A recovery catalog for prod1 is only effective if it is separated from the data that it is designed to protect.

An Oracle customer, who is an ASP provider, manages 200+ databases. He created one catalog to keep backup information for all of the databases. The catalog provides a centralized repository of backup history. The DBA wrote a report to summarize the backup status of all databases by querying the catalog tables.

Merrill Lynch manages over 145 production and test databases. They have created two separate RMAN catalogs in the same database. This can be achieved by creating two schemas into which the RMAN catalog is created. One catalog contains the production backup history while the other maintains the test database backup information.

The recovery catalog need not be available to perform a database backup. The target's control file records RMAN backup information. When the recovery catalog is used again, the control file is used to synchronize the data in the catalog.

### ***BACK IT UP***

How do you backup the Recovery Manager catalog? You can use the control file of the catalog database as the RMAN repository. While using RMAN as the backup tool for the catalog, the control file contains the backup information in case a restore is required. A logical backup of the catalog can be made using Oracle's export utility. Replication of the catalog may be achieved by setting up a physical standby database on another host.

#### **Top 10 Reasons to integrate Recovery Manager into your Backup and Recovery Strategy**

- |   |   |
|---|---|
| <b>10.</b> Extensive Reporting            | <b>5.</b> Easily integrates with Media Managers |
| <b>9.</b> Incremental Backups             | <b>4.</b> Block Media Recovery (BMR)            |
| <b>8.</b> Downtime Free Backups           | <b>3.</b> RMAN knows Archive Logs               |
| <b>7.</b> Backup and Restore Validation   | <b>2.</b> Corrupt Block Detection               |
| <b>6.</b> Backup and Restore Optimization | <b>1.</b> Trouble Free Backup and Recovery      |

## **COMING ATTRACTIONS FOR BACKUP AND RECOVERY**

RMAN continues to add features and improve capabilities so that recovery of the enterprise is fast and efficient. In the next release of Oracle9i, backup and recovery will add the capabilities to backup the SPFILE, throttle the amount of archivelogs restored during a database recovery, and offer an advisory of bounded instance recovery.

### ***SPFILE BACKUP***

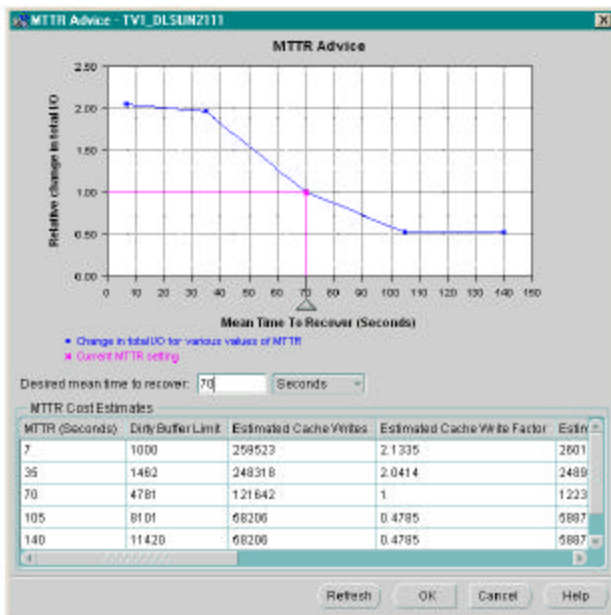
RMAN will automatically back up the current SPFILE. When included in a backup, SPFILE will be backed up together with the control file. It will be possible to restore any SPFILE file backup to an arbitrary location. If a restore location is not specified, then RMAN will restore SPFILE to the location from which the SPFILE was backed up.

### RESTORE ARCHIVE LOG THRESHOLD

RMAN can control the amount of archived logs restored during recovery, so they do not exceed a certain amount of space. By setting the parameter, you can specify of disk space for restored archived logs. This prevents an out-of-disk space condition during a critical database recovery.

### MEAN-TIME-TO-RECOVER ADVISORY

In Oracle9i, initialization parameter `FAST_START_MTTR_TARGET` (MTTR) is used to specify the amount of time Oracle is expected to perform recovery should a crash occur. Based on its value, the system statistics gathered from past recoveries and other system parameters, Oracle automatically adjusts the checkpointing behavior to meet the expected recovery time. In general, the smaller the MTTR setting is, the more aggressive the system performs checkpointing. The smaller the MTTR, the more total IO the system would have.



There is a trade-off between recovery time and run-time operation performance. If you desire a shorter recovery time, you may have to pay the cost of more IO during run-time. The MTTR advisory provides you with an advisory on how different values of `FAST_START_MTTR_TARGET` affect the number of IO in the system. After the system runs a typical workload for a while, a re-query of the advisory estimates the number of cache writes under other MTTR settings. By looking at the different MTTR settings and their corresponding cache write ratio, you can decide which MTTR value fits the user's recovery and database performance needs.

### SUMMARY

With the increasing database sizes and the demand for continuous database server availability, the time frame available for doing backup and recovery is becoming smaller and smaller. Hence the option of shutting down the database for an offline full backup is less available even in non mission-critical environments. Achieve your enterprise MTTR by reviewing and testing your recovery plan. Incorporating efficient backup tools and methods ensure the recoverability of your database is possible. The key to recovery is the backup.