

# Real-Life Recovery: Perspective, Preparation and Performance

Daniel W. Fink  
7S Consulting, Inc.

In an anonymous office building downtown, an Oracle database administrator checks the log files to verify that the backup processes ran properly last night. Then he grabs a 2<sup>nd</sup> cup of coffee and moves on to the day's tasks, emergencies and assorted events. Users, managers and developers are concerned with data accuracy, proper security and blinding performance, not in the mundane task of testing recovery. Little does he or the company know that the past month's backups are useless. Since the recent operating system update, the tape drive, the only tape drive on site, writes without error, but cannot read a single bit that is on the tape.

A few miles away, 2 DBAs and a System Administrator are completing a 12-hour system and database recovery. No one has slept in over 36 hours. Mistakes were made, but they were caught well before the right pinkie was poised over the <ENTER> key. When the final command was issued, the three felt suddenly refreshed...success is a great stimulant. For the 3<sup>rd</sup> time in as many months, they have successfully performed a recovery with no perceivable loss of data. The Operations Manager walks in with a pot of coffee and a piece of paper. "Pretty good, no data loss," she says, "but you could have been done 2 hours ago. I'll grade it a B+. Let's go for an A next time! Go home, get some rest and we'll see you tomorrow."

These two tales illustrate several important points regarding Oracle Backup and Recovery. Most sites focus on backups and other issues before addressing recovery. Single points of failure are shortcuts to disaster, both business and career. The only measure of a successful backup plan is how completely recovery can be performed. The time to practice a recovery is not when a production system is down and the Order Entry clerks are waiting. Personnel can be the single greatest weakness in recovering a database, but they can also be the single greatest strength.

## ➤ Perspective

Most Oracle documentation gives equal billing to backup and recovery. The routine task of setting up, monitoring and maintaining backup processes consumes valuable resource time. Once the process has been established, other issues, such as security, tuning and troubleshooting move to the top of the DBA task list. Usually, a crisis is the time that the recovery process is performed. That is also the worst possible time to find out that the backup process is flawed, the SYSTEM tablespace file cannot be read from the tape.

This situation is a result of the attitude that the DBA is responsible for performing database backup. **The actual responsibility is to restore or recover the database to the point in time and within the downtime window determined by the business needs.**

Rather than the backup process dictating the recovery options, the recovery requirements should drive the backup process. A backup strategy is determined by the answer to a single question, "What is the best method of backup and archiving that meets all of the business needs of data loss, downtime and cost?" As with many decisions, the question appears to be simple, but is complex when defining 'best method', 'business needs', 'data loss', 'downtime' and 'cost'. In an attempt to define these variables, the Cost v. Lost Revenue Model is used to assist in answering this question.

## • Restore, Recover, Rebuild

To restore is to return to a former state. In database terms, this means to return the database file(s) to a previous point in time determined by copying the file from backup to online media. The database or particular file may be missing transactions. Restoration is required before database recovery can be performed after a media or other similar failures. If the database is not

running in archivelog mode, this is probably the opened state of the database.

To recover is to return to a normal state. In database terms, this means to return the database to a point in time determined by the situation. It may be a complete or incomplete recovery, depending upon the circumstances. Files are restored, and then archived redo logs are applied to bring the database to the desired state. Some transactions may be lost, but they are usually few.

To rebuild is to recreate the state. In database terms, this usually involves reusing the data and process that populated the database in the most recent iteration. In some environments, particularly data warehouse or nonproduction, the raw data exists outside of the database. There may be no actual transactions since the last load therefore nothing is lost.

## • **Base Knowledge and Training**

Transaction Architecture. The reason for a database is to support transactions, whether they are entering a \$.99 cable order or running a 3-day Regional Sales data warehouse query. The understanding of how a transactions works, and what happens when it does not work, are the cornerstone to database knowledge.

Recovery Structures and Processes. A DBA must be able to answer the question "What is the difference between a rollback segment and redo log?" in 25 words or less. Know the recovery structures, what they do and how to recover from the loss of each one. It is important to know the performance implications of running a database in archivelog mode, both from a processing standpoint and archivelog\_destination cleanup. Intimate knowledge of each recovery type is critical to success. Minimum areas of understanding include the pros and cons of each type, impact on business and downtime and required steps, i.e. which files to restore, what steps to perform during and after recovery, etc.

Business Requirements. Each database is different. Hours of operation, failure tolerance, data integrity rules, etc. A database is a dynamic system, constantly changing within its environment, which is also changing. The requirements in place 6 months ago may have changed drastically. Define your recovery goals

for each system and database. Decisions are driven by the amount of time allowable to perform a recovery and the amount of data allowed to be lost. These decisions are not set in concrete; they will change over the lifetime of the database.

Training and Education. The responsibility for performing backup and recovery operations should never fall on the shoulders of a person who is not well trained and educated in the ways of Oracle. Backup development and practice recovery can be done in a mentor-student relationship, until the student is confident in the tasks to attempt perform these unassisted. The process is similar to learning to fly a plane, put in the hours before attempting to go solo. There are several excellent books and white papers available regarding this discussion. One of the best is Rama Velpur's *Oracle Backup & Recovery Handbook* from Oracle Press.

## • **Points of Failure**

'Never have a single point of failure' is the Oracle Backup & Recovery Maxim. The basic foundation is that a single failure should never prevent a backup or recovery process from succeeding. If day-old archive logs are deleted after the cold backup, the backup tape is a single point of failure. If the logfiles are unable to be restored, the previous backup cannot be used to recover the database because the archived logs are unavailable.

Identify the 'show stoppers' and have at least 2 options for bypassing. A single bad backup or command should never be the cause of a failed recovery. Taking an extra 2 hours to recover a lost datafile from a 2-day-old backup is much easier than explaining why yesterday's sales were lost because the archived redo log was lost.

Prevention of failures is a proactive approach to recovery. Recovery from a dropped table is possible, but implementing and enforcing database security, access and maintenance procedures can prevent the situation. Using an appropriate RAID configuration can minimize the impact of disk loss. There are expenses associated with prevention, but they may be minimal compared with actual recovery costs. Even if proactive solutions are implemented, there will

still be areas of weakness. While certain RAID configurations will protect against the loss of a single disk, they cannot protect against the delete command.

- **Failure Prone Components**

Personnel The most failure prone component is the person issuing the commands. Document the steps required for each recovery type. When documenting, assume little, if any practice. The recovery manual should address the areas that pose the greatest chance of total loss of service. Use signposts like “Proceed with Caution”, “Stop & Review”, etc. and decision trees. Update the manual periodically, at least at every major change. If you are using Oracle Support assistance, make certain that you are talking to someone who has performed recoveries.

Redo Logs (online, offline and archived) The loss of a single redo log can result in total recovery failure. This is one reason that Oracle recommends using mirrored redo log groups. Using hardware level mirroring is not sufficient, you are protected from media failure, i.e. disk loss, but you are not protected from the delete command. Archived logs should exist in their natural state on at least one backup tape. Once there, they may be compressed and kept online for several days, and placed on several other backup tapes. This strategy offers good protection from loss and assists in recovery speed by having the logs online. Although rare, the compression/uncompression cycle may result in corrupted files, so use with caution. If disk space allows, have at least two uncompressed archive log copies on different tapes.

Control Files If all control files are lost, Oracle offers several methods for performing recovery. However, these tasks are risky and easily avoidable. Another Oracle maxim is to have 3 copies of your controlfile, because a corrupt controlfile will always, well almost always, be copied over a good controlfile. If you have a 3<sup>rd</sup> controlfile, the erroneous copy will still happen, but only you will know...the 3<sup>rd</sup> copy is used to save the day and your job.

Initialization Parameter Files These files are often overlooked. As with controlfiles, their loss can be overcome, but at a substantial effort. Data is not

lost, but downtime is increased. A copy of all parameter files and a dump of v\$parameter add several layers of backup.

System Tablespace datafile(s) These datafiles may be the most important in the database for it is the blueprint. Without the data dictionary, Oracle is blind to all of the data that is residing in the datafiles.

Data Tablespace datafile(s) Under the OFA standard, data and indexes are separated into distinct tablespaces. This allows backup and recovery plans to be flexible. If the system and data tablespace datafiles are recoverable, all other tablespaces can be built from scratch, if sufficient documentation exists. For example, indexes can be recreated, if the most current DDL scripts are available. For Oracle8, data tablespaces include those objects defined as index-only tables.

Rollback Segment Tablespace datafile(s) Although these tablespaces are used by transactions to provide transaction recovery and read consistency, the loss of this tablespace is not necessarily a ‘death sentence’. The key factor is the state of the database at the point of the backup. If it is consistent mode, i.e. no uncommitted transactions, there should be no rollback entries. As such, the segments can be dropped and rebuilt without data loss.

Index Tablespace datafile(s) As stated above, if the most current DDL scripts exist, indexes can be recreated from existing data. There is a substantial recovery performance penalty incurred as the indexes are being rebuilt, but no data loss occurs.

Temporary Tablespace datafile(s) Once the database is shutdown, the temporary segments are ‘clean’. If the temporary tablespace(s) are lost, recreation is a fairly simple process. As with indexes, there is a recovery performance penalty, but no data loss occurs.

Misc Files File maps, object recreation scripts, backup scripts and other assorted files should be part of the backup process. Performing a media recovery is greatly simplified if an up-to-date tablespace-file-device map is available. While these files are not required for recovery, they can shorten and simplify the recovery process.

## ➤ Preparation – the Cost v. Lost Revenue(CLR) Model

In a perfect world, the CIO would hand the DBA and SA a blank check for backup hardware/software and provide an unlimited operations budget. In the real world, there is a balance between the costs associated with backup and the benefits of recovery. Although the business needs drive the Cost-Benefit Analysis, it is the responsibility of the technical staff to educate the users and support the decision-making process. It is not the responsibility of the technical staff to make the decision.

The basic concept is to compare the cost of backup and recovery with lost revenue. This is balanced so that the frequency of backup v. the frequency of recovery is appropriately weighted. In a production environment, a downed Order Entry system can be quantified. In a development environment, the quantification is related more to lost Developer/User time.

The model should be generated for 3 scenarios – worst, best and anticipated. The anticipated case should be somewhere between worst and best. Best case should not be a fantasy, assume some level of recovery requirements.

There are five basic steps to the CLR Model process:

### 1) Educate the decision makers

Discuss issues in clear, nontechnical terms that are clear to everyone. The first step is to lay the foundation by educating the decision-makers in basic knowledge of backup and recovery. Although an Accounting Manager does not need to understand the intimate details of archive logging, they do need to understand that transaction recovery is not possible without it...and what the business implications are of adopting this strategy.

### 2) Assign costs to each resource.

This phase is primarily devoted to assigning actual or opportunity costs to each resource. The resources are hardware, software and peopleware. The backup process will consume disk space and/or tapes, CPU cycles and memory on the host platform. Backup software must be

purchased or written. Administrators of various experience levels must monitor backups, perform tests, document and perform the occasional recovery.

For hardware and software, acquisition and operation costs can be quantified or estimated. The operation costs are defined within the scope of time. The actual time is not important, but it must be consistent. Hardware operation costs for a year are not comparable to backup operation costs for 6 months.

For personnel, the actual or opportunity cost is determined. For most production environments, personnel cost is actual, i.e. the sum of salary and benefits or hourly charges. For development environments, the cost may be classified as opportunity, i.e. the hourly rate that the client is being charged.

Although they are not assigned at this point, the intangible costs should be discussed. These costs include impact on project timelines, goodwill with users, management and technical personnel. If a project is approaching a critical junction, the cost of downtime may escalate dramatically. Users may tire of experiencing excessive downtime; Managers may tire of user complaints; Developers may tire of reloading the past month's data; Administrators may tire of 36-hour days. The short-term impact may be less money spent on the systems, but the long-term impact may be frustration, high turnover and added costs.

### 3) Assign costs to each backup strategy

There are three types of backup strategy: hot, cold with archiving and cold without archiving. A cost for each scenario is calculated in two areas: fixed and variable.

Fixed costs are the one-time costs for each step during the operational period. Software needs to be purchased or written once. Disk space required for Archived Logs or backup datafiles is another one-time investment. These costs are incurred regardless of the frequency of backups, assuming the frequency is greater than 0.

Variable costs are incurred each time the backup process occurs. Logfiles must be monitored, tapes must be written and stored and scripts must be maintained.

At this point, the first business decision is made. The anticipated frequency of the database backups is determined. This decision is very preliminary, it is not set in stone. The less frequent the backup, the more costly the recovery, but a final decision is premature. It is important to use the preliminary decision as a baseline. Once the recovery costs are determined, the type of backup can be revisited.

Another alternative is to complete this phase of the model for each type and frequency of backup. While time consuming, it may be simpler and quicker when busy schedules and diverse audiences are involved.

#### **4) Assign costs to each recovery scenario**

The most common types of recoveries are object, file and disk, each with a distinct strategy. Cost determination is common among all of the strategies, and is slightly different from determining backup cost.

The first area of recovery cost is practice cost. The major fixed expense is for a practice platform. An ideal situation is to have a 'sandbox' for System and Database administrators to test upgrades and recoveries. In many situations, an older development platform may be adequate. Added to the cost is the variable personnel cost for each practice.

The second area of cost is restoration. This includes the time to detect and correct/bypass the cause of failure. Then the backup media is retrieved and objects or files are restored. If archiving is being used, only the affected object or file needs restoration, which can significantly reduce the restoration time.

A final, but optional, area is the cost of recovery. If archiving is being used, Oracle recovery processes can be used. This is usually the quickest method of recovery. Other methods are transaction reentry and data reload, which are both time and personnel resource intensive. For a small decision support system, it may be more efficient to recreate the database from tested scripts and reload the previous period's data.

The best method for estimating recovery time is to actually practice each step. An alternative method

is to use estimations for each step. Regardless, an qualified guess is better than nothing, but remember that these decisions are among the most important that may be made in your career.

#### **5) Support the business user(s) in making an educated decision.**

Once all of the costs are determined, the process of finding an acceptable balance is begun. Unless the company can afford hot-standby, fault-tolerant, fail-over systems, compromises must be made. Simply put, it is time to make an educated, documented decision that is supported by the business. The decisions that are made during this step may result in lost data, missed sales, overtime and tough explanations to management. If it is determined that the decision requires additional resource, such as disk drives or a tape library, they can be purchased before they are needed. Recovery planning is part of the system design phase.

Although it may not be easy explaining to management to invest \$100,000 in a system that is rarely used, it is better than explaining why \$1,000,000 in sales were lost yesterday because the Order Entry system is not archiving transactions.

#### **➤ Performance**

Oracle database recovery is one area where the line between success and failure is clear. If the database is recovered according to the defined business requirements, the recovery was successful, if not, the recovery was a failure. Trying to successfully recovery is not acceptable. In the words of Yoda, "Do or Do Not. There is no Try."

The only valid method for testing a backup process is to perform a test recovery. This test should be repeated on a regularly scheduled basis and after any major change, such as a new tape device or system upgrade. The time to find out that the tape drive is not functioning properly or that a new database file is not being backedup is during a practice run, not a live recovery situation. Practice runs will also expose gaps in procedures, documentation, training and, most importantly, confidence.

- **Practice, Practice, Practice**

Practice recoveries are the only method for gaining confidence in the process and exposing weaknesses. Being able to quote chapter and verse from recovery theory is not a substitute for drawing upon actual experience, especially at 4:00 in the morning after a 20-hour day. At that time, you may be one keystroke away from success or failure. It is also the time when your mental faculties are at their most vulnerable.

Mistakes during a practice recovery cause no harm. In fact, they are a tremendous learning experience. If you do not have the time and training to perform a recovery 100% right the first time, there may be no next time...at least not at the current company.

Practice will expose the 'bridge burning' steps. These points are critical decisions where returning to a previous step is difficult or impossible, unless certain precautions are taken. Copying a backup datafile in place of an existing datafile is such a step. These steps must be documented and well understood. It may be appropriate to backup a 'bad' datafile prior to restoration. This allows for a retreat should it be needed. Treat these situations as 'show stoppers'.

Oracle8 introduced the Recovery Manager tool, which joins a number of 3<sup>rd</sup> party products. If the business uses a tool, it needs to be part of the practice. You must also plan for the possibility that the tool is unavailable, always practice manual forms of recovery. Tool availability is not a substitute for basic foundational knowledge.

Practicing recovery is not a minor expense, but it pales in comparison to losing a substantial amount of data for a production environment. It requires professional dedication to the business. As recoveries are practiced, knowledge and skill are developed and confidence is increased. In turn, this may assist in redeploying the backup and recovery strategy.

- **Executing the Plan**

1. Stop panicking. It is imperative to get into the proper frame of mind. Taking an extra 15 minutes to calmly discuss the possible alternatives may make the difference in a

failed recovery attempt and an additional 15 hours of downtime.

2. Determine scope and cause of failure. The type of failure will have serious impact on the types of recovery to be considered. If the failure will be repeated shortly after the recovery is performed, the recovery is wasted. If a table has been dropped, performing a complete recovery will return the database to the point AFTER the table was dropped.
3. Correct/Bypass failure. When a disk fails, the datafiles should be restored to another disk or the disk should be replaced. If the failure is not corrected or bypassed, a recovery may be wasted.
4. Identify Plan of Attack. Once the failure is understood and corrected, the recovery options are defined and a plan is determined. If a datafile for an index-only tablespace has been damaged, it may be more efficient to drop the tablespace and rebuild the affected indexes. Depending upon the plan, the next step may not be required.
5. Restore affected data. After failure correction, the process of restoration can begin. A tape or other backup media is retrieved and the affected data and log files are restored to the system. If redo log files are not being archived, the complete database must be restored.
6. Perform recovery. The actual recovery strategy that was determined in the plan is performed. In the above example, this may involve dropping the tablespace, recreating it and then rebuilding the indexes and recreating constraints. Depending upon the backup plan, there may be substantial resource requirements, especially if the data must be reloaded or transactions reentered. If time allows, a full cold backup should be performed after recovery but before the database is opened for general use.
7. Postmortem. Debrief, document and determine improvements. Determine if the failure situation could be prevented. Losing data and incurring downtime due to a preventable failure is rarely acceptable in a business critical system, unless the business

requirements have accounted for and accepted this possibility.

- **Conclusion**

Among the myriad of DBA tasks and responsibilities, the most important to the business is the ability to properly recover from a failure. There are many factors that determine the type and scope of recovery, primarily the balance between what the business can afford to spend and what it can afford to lose. These decisions are not the responsibility of the DBA or SA; rather, they are to be made by key members of the user community. The technical staff functions as a support organization by educating the users and assisting in the decision-making process using a Cost v. Lost Revenue Model.

Once the business has determined the backup and recovery needs, the technical staff becomes responsible for insuring that these operations are properly executed. As a technical administrator, you are the greatest strength and greatest weakness for the recovery process. Training and practice are the paths to success. The recovery process is too critical to depend on less than 100% effort and ability.