
ORACLE®

ORACLE®

RMAN Configuration and Performance Tuning Best Practices

Timothy Chien
Principal Product Manager
Oracle Database High Availability
Timothy.Chien@oracle.com

Agenda

- Recovery Manager Overview
- Configuration Best Practices
  - Backup Strategies Comparison
  - Fast Recovery Area (FRA)
- Performance Tuning Methodology
  - Backup Data Flow
  - Tuning Principles
  - Diagnosing Performance Bottlenecks
- Summary/Q&A

Oracle Recovery Manager (RMAN)

Oracle-Integrated Backup & Recovery Engine

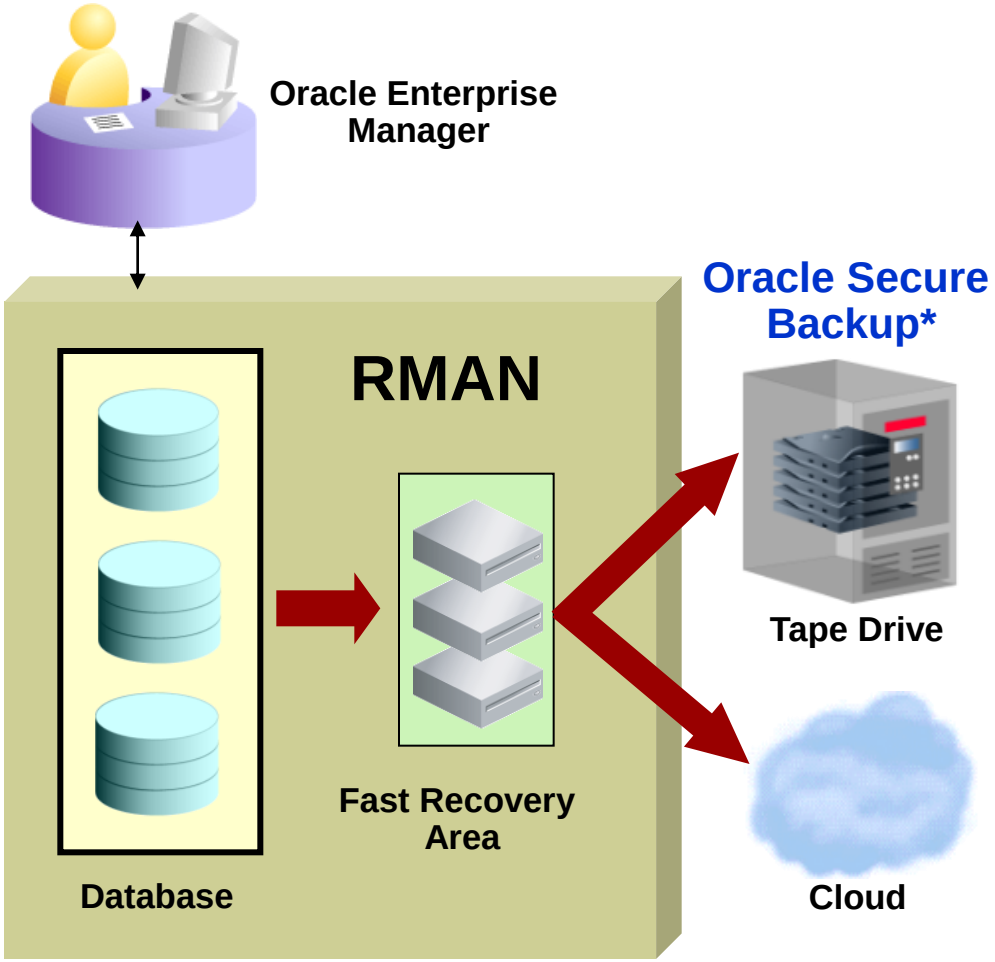

- Intrinsic knowledge of database file formats and recovery procedures
  - Block validation
  - Online block-level recovery
  - Tablespace/data file recovery
  - Online, multi-streamed backup
  - Unused block compression
  - Native encryption
- Integrated disk, tape & cloud backup leveraging the Fast Recovery Area (FRA) and Oracle Secure Backup

***RMAN also supports leading 3rd party media managers**

Critical Question to Ask

- What are my recovery requirements?
  - Assess tolerance for data loss - *Recovery Point Objective (RPO*)
    - How frequently should backups be taken?
    - Is point-in-time recovery required?
  - Assess tolerance for downtime - *Recovery Time Objective (RTO)*
    - Downtime: Problem identification + recovery planning + systems recovery
    - Tiered RTO: database, tablespace, table, row
  - Determine backup retention policy
    - Onsite, offsite, long-term
- Then..how does my RMAN strategy meet those requirements?

Option 1: Full & Incremental

Tape Backups

- **Well-suited for:**
  - Databases that can tolerate hours/days RTO
  - Environments where disk is premium
  - Low-medium change frequency between backups, e.g. < 20%
- **Backup strategy:**
  - Weekly level 0 and daily ‘differential’ incremental backup sets to tape, with optional backup compression
  - Enable block change tracking - only changed blocks are read and written during incremental backup
  - Archived logs are backed up and retained on-disk, as needed

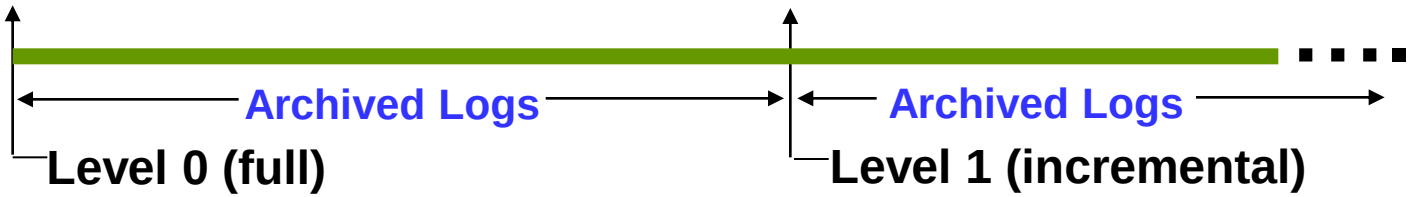

RMAN Script Example

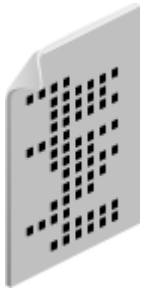

- Configure `SBT` (i.e. tape) channels:
  - **`CONFIGURE CHANNEL DEVICE TYPE SBT PARMS '<channel parameters>';`**
- Weekly full backup:
  - **`BACKUP AS BACKUPSET INCREMENTAL LEVEL 0 DATABASE PLUS ARCHIVELOG;`**
- Daily incremental backup:
  - **`BACKUP AS BACKUPSET INCREMENTAL LEVEL 1 DATABASE PLUS ARCHIVELOG;`**

Option 2: Incrementally Updated
Disk Backups

- **Well-suited for:**
  - Databases that can tolerate no more than a few hours RTO
  - Environments where disk can be allocated for 1X size of database or most critical tablespaces
- **Backup strategy:**
  - Initial image copy to FRA, followed by daily incremental backups
  - Roll forward copy with incremental, to produce new on-disk copy
  - Full backup archived to tape, as needed
  - Archived logs are backed up and retained on-disk, as needed
  - Fast recovery from disk or **`SWITCH`** to use image copies

RMAN Script Example

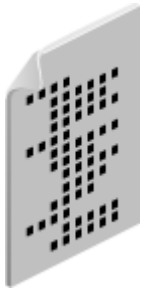

- Configure SBT channels, if needed:
  - `[CONFIGURE CHANNEL DEVICE TYPE SBT PARMS '<channel parameters>';]`
- Daily roll forward copy and incremental backup:
  - `RECOVER COPY OF DATABASE WITH TAG 'OSS';`
  - `BACKUP DEVICE TYPE DISK INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG 'OSS' DATABASE;`
  - `[BACKUP DEVICE TYPE SBT ARCHIVELOG ALL;]`
- What happens?
  - First run: Image copy
  - Second run: Incremental backup
  - Third run+: Roll forward copy & create new incremental backup
- Backup FRA to tape, if needed:
  - `[BACKUP RECOVERY AREA;]`

Fast Recovery with RMAN `SWITCH` Demo

Option 3: Use Data Guard

Offload Backups to Physical Standby

- **Well-suited for:**
  - Databases that require no more than several minutes of recovery time, in event of any failure
  - Environments that can preferably allocate symmetric hardware and storage for physical standby database
  - Environments whose tape infrastructure can be shared between primary and standby database sites
- **Backup strategy:**
  - Full & incremental backups offloaded to physical standby database
  - Fast incremental backup on standby with Active Data Guard
  - Backups can be restored to primary or standby database
    - Backups can be taken at each database for optimal local protection

Comparison: Backup Strategies

Strategy	Backup Factors	Recovery Factors
Option 1: Full & Incremental Tape Backups	• Fast incrementals • Save space with backup compression • Cost-effective tape storage	• Full backup restored first, then incrementals & archived logs • Tape backups read sequentially
Option 2: Incrementally Updated Disk Backups	• Incremental + roll forward to create up-to-date copy • Requires 1X production storage for copy • Optional tape storage	• Backups read via random access • Restore-free recovery with **`SWITCH`** command
Option 3: Offload Backups to Physical Standby Database	• Above benefits + primary database free to handle more workloads • Requires 1X production hardware and storage for standby database	• Fast failover to standby database in event of any failure • Backups are last resort, in event of double site failure

Fast Recovery Area (FRA) Sizing

- If you want to keep:
  - Control file backups and archived logs
    - Estimate total size of all archived logs generated between successive backups on the busiest days x 2 (in case of unexpected redo spikes)
  - Flashback logs
    - Add in {Redo rate x Flashback retention target time x 2}
  - Incremental backups
    - Add in their estimated sizes
  - On-disk image copy
    - Add in size of the database minus size of temporary files
  - Further details:
    - http://download.oracle.com/docs/cd/E11882_01/backup.112/e10642/rcmconfb.htm#i1019211

FRA File Retention / Deletion Policies

- When FRA space needs exceed quota, automatic file deletion occurs in the following order:
  1. Flashback logs
     - Oldest Flashback time can be affected (with exception of guaranteed restore points)
  2. RMAN backup pieces/copies and archived redo logs that are:
     - Not needed to maintain RMAN retention policy, or
     - Have been backed up to tape (via **`DEVICE TYPE SBT`**) or secondary disk location (via **`BACKUP RECOVERY AREA TO DESTINATION '..'`**)
- If archived log deletion policy is configured as:
  - **`APPLIED ON [ALL] STANDBY`**
    - Archived log must have been applied to mandatory or all standby databases
  - **`SHIPPED TO [ALL] STANDBY`**
    - Archived log must have been transferred to mandatory or all standby databases
  - **`BACKED UP <N> TIMES TO DEVICE TYPE [DISK | SBT]`**
    - Archived log must have been backed up at least **`<N>`** times
  - If `[`**`APPLIED`** or **`SHIPPED`**`]` and **`BACKED UP`** policies are configured, both conditions must be satisfied for an archived log to be considered for deletion.

Performance Tuning Methodology

- RMAN Backup Data Flow
- Performance Tuning Principles
- Diagnosing Performance Bottlenecks

RMAN Backup Data Flow

A. Prepare backup tasks & read blocks into input buffers
B. Validate blocks & copy them to output buffers
   - Compress and/or encrypt data if requested
C. Write output buffers to storage media (**`DISK`** or **`SBT`**)
   - Media manager handles writing of output buffers to **`SBT`**

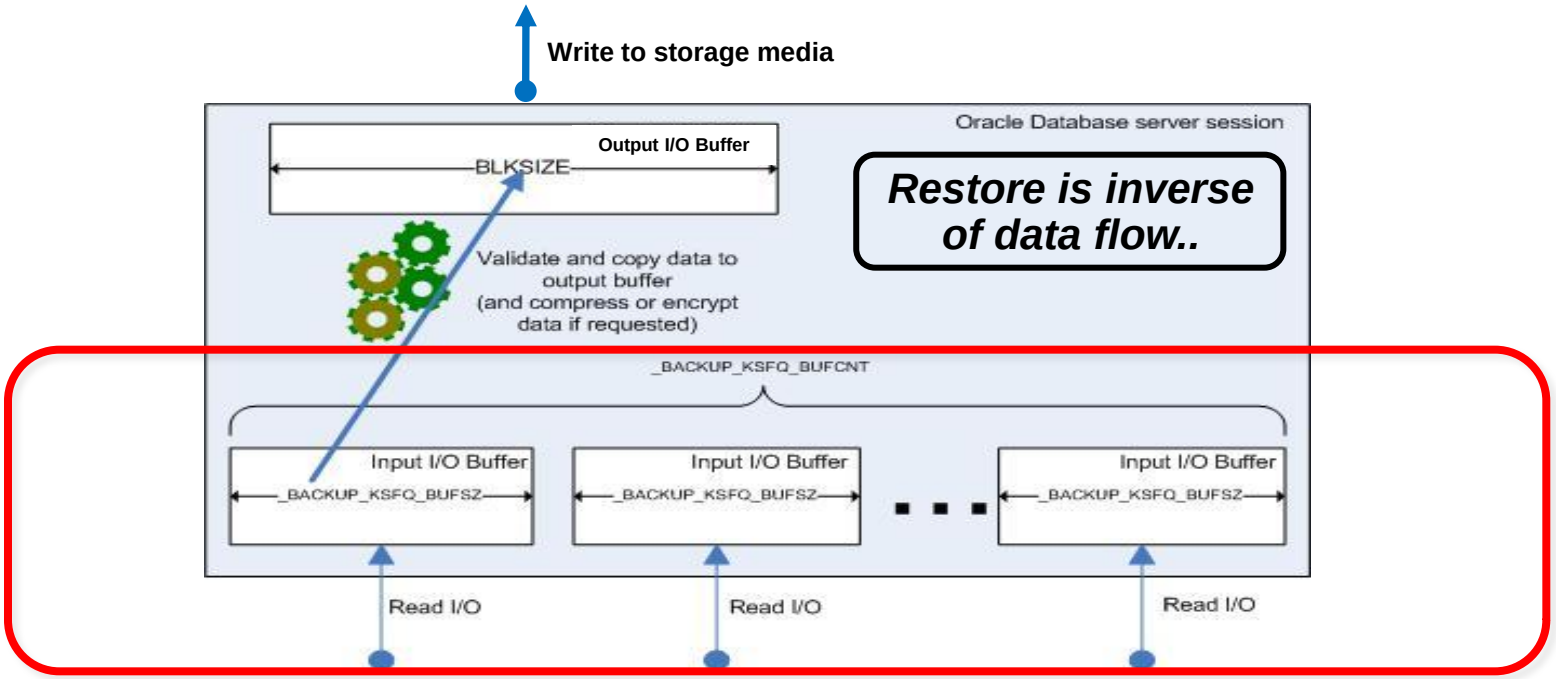

Tuning Principles

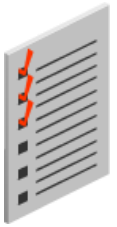

1. Determine the maximum input disk, output media, and network throughput
   - E.g. Oracle ORION
     - http://download.oracle.com/docs/cd/E11882_01/server.112/e16638/iodesign.htm#CACJEEDI
   - Evaluate network throughput at all touch points, e.g. database server -> media management environment -> tape system (ref. TCP/IP performance measurement tools such as **`qperf`**)

2. Configure disk subsystem for optimal performance
   - Use ASM: typically, separate DATA and FRA disks
   - If not using ASM, stripe data files across all disks with 1 MB stripe size

Tuning Principles … contd.

3. Tune RMAN to fully utilize disk subsystem and tape
   - Verify asynchronous I/O supported by platform, otherwise:
     - Disk backup: set **DBWR_IO_SLAVES**
     - Tape backup: set **BACKUP_TAPE_IO_SLAVES** unless media manager states otherwise
   - Disk backup: allocate as many channels as can be handled by system
     - For image copies, one channel processes one data file at a time
   - Tape backups: allocate one channel per tape drive
     - Note: Restore time will degrade with higher number of channels per tape drive, due to tape-side multiplexing (which leads to interleaving of disjoint backup sets in the same tape)
     - With higher number of channels, if read phase time (determined by **BACKUP VALIDATE**)
       - Decreases (vs. same number of channels), bottleneck is in read phase: ref. next slide
       - Stays the same, bottleneck is most likely in media manager: ref. Slide 27

Read Phase Tuning

RMAN Multiplexing

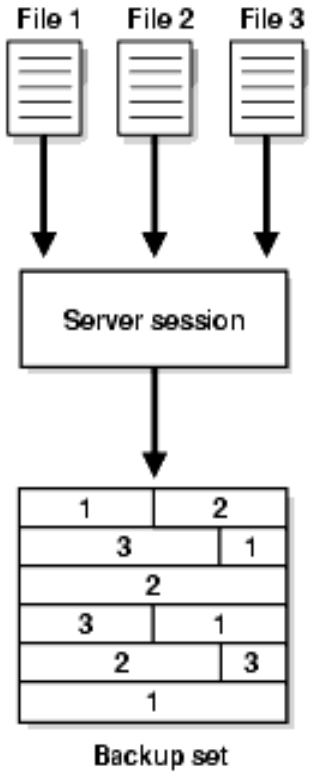

- Multiplexing level: max number of files read by one channel, during backup - controlled by:
  - **FILESPERSET** [default: 64] parameter of the **BACKUP** command: how many datafiles to put in each backup set
  - **MAXOPENFILES** [default: 8] parameter of **ALLOCATE CHANNEL** or **CONFIGURE CHANNEL**: how many datafiles RMAN can read from simultaneously.
- Multiplexing level: Min (**MAXOPENFILES**, **FILESPERSET**)
  - Determines number and size of input buffers in **V$BACKUP_ASYNC_IO/V$BACKUP_SYNC_IO**
  - All buffers allocated from PGA, unless disk or tape I/O slaves are enabled
    - If slaves are enabled, all buffers allocated from SGA or **LARGE_POOL** (if set)

Read Phase Tuning … contd.

RMAN Input Buffers

- For ASM or striped system:
  - **`MAXOPENFILES=1`** => 16 buffers/file, 1 MB/buffer = 16 MB/file
    - Allows largest number and sized buffers by default
    - Additional multiplexing not needed, since files are striped
- For non-striped system:
  - **`MAXOPENFILES=8`** => 4 buffers/file, 512 KB/buffer = 2 MB/file
    - Reduce the number of input buffers/file to more effectively spread out I/O usage (since each file resides on one disk)

Tuning Principles … contd.

4. If **`BACKUP VALIDATE`** still does not utilize available disk I/O & there is available CPU and memory:
   - Increase RMAN buffer memory usage
     - With Oracle Database 11*g* Release 11.1.0.7 or lower versions -
       - Set **`_BACKUP_KSFQ_BUFCNT`** (default 16) = # of input disks
         - Number of input buffers per file
         - Achieve balance between memory usage and I/O
       - Set **`_BACKUP_KSFQ_BUFSZ`** (default 1048576) = stripe size (in bytes)
     - With Oracle Database 11*g* Release 2 -
       - Set **`_BACKUP_FILE_BUFCNT,_BACKUP_FILE_BUFSZ`**
     - Restore performance can increase with setting these parameters, as output buffers used during restore will also increase correspondingly
     - Refer to ***Support Note 1072545.1*** for more details
     - Note: With Oracle Database 11g Release 2 & ASM, all buffers are automatically sized for optimal performance

RMAN Backup Data Flow

A. Prepare backup tasks & read blocks into input buffers

B. Validate blocks & copy them to output buffers
   - Compress and/or encrypt data if requested

C. Write output buffers to storage media (`DISK` or `SBT`)
   - Media manager handles writing of output buffers to `SBT`

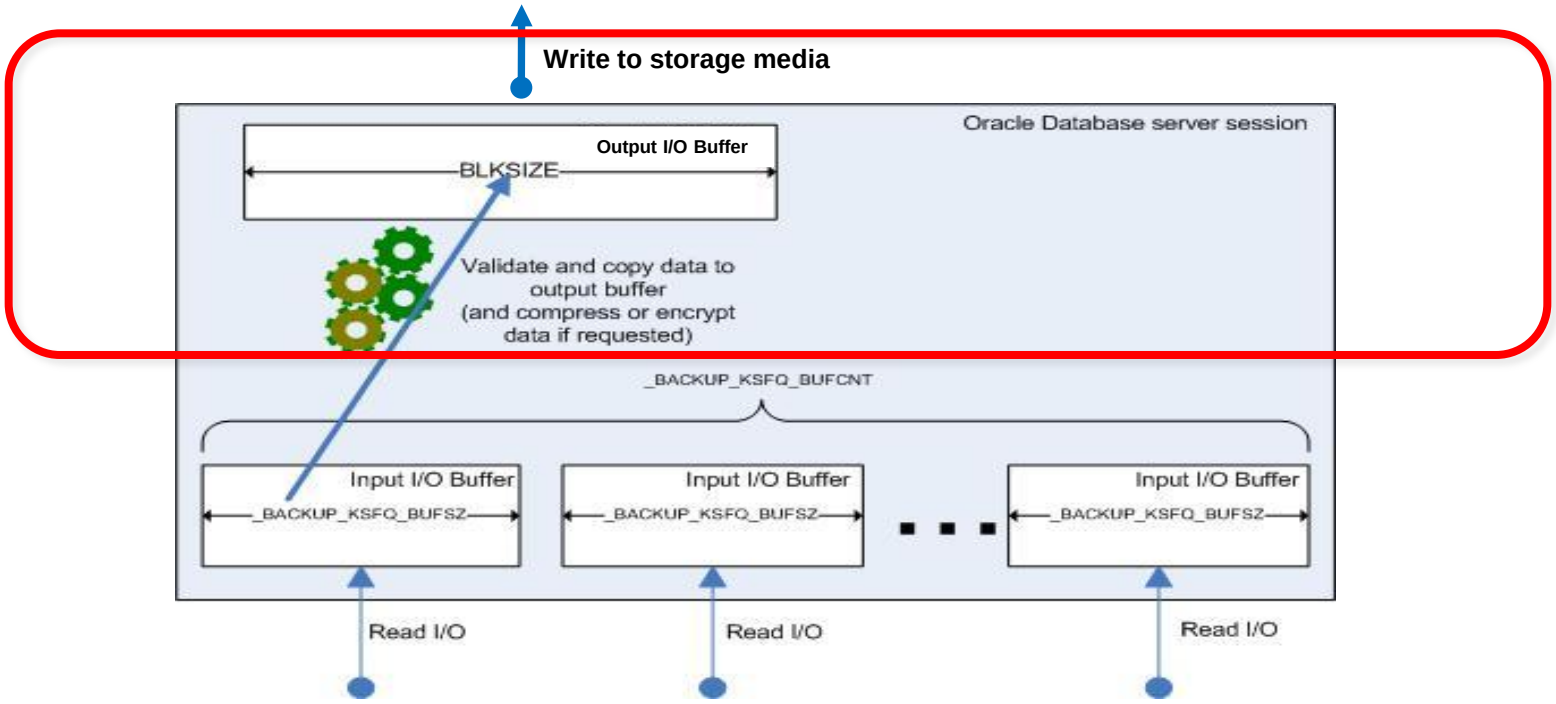

Tuning Principles … contd.

5. RMAN backup compression & encryption guidelines
   - Both operations depend heavily on CPU resources
   - Increase CPU resources or use **`LOW/MEDIUM`** setting
   - Verify that uncompressed backup performance scales properly, as channels are added
   - For encryption:
     - TDE column encryption
       - For encrypted backup, data is double encrypted (i.e. encrypted columns treated as if they were not encrypted)
     - TDE tablespace encryption
       - For compressed & encrypted backup, encrypted tablespaces are decrypted, compressed, then re-encrypted
       - If only encrypted backup, encrypted blocks pass through backup unchanged

Tuning Principles … contd.

6. Tune RMAN output buffer size
   - Output buffers => blocks written to **`DISK`** as copies or backup pieces or to **`SBT`** as backup pieces
   - Four buffers allocated per channel
   - Default buffer sizes
     - **`DISK: 1 MB`**
     - **`SBT: 256 KB`**
       - Adjust with **`BLKSIZE`** channel parameter
       - Set **`BLKSIZE`** >= media management client buffer size
       - No changes needed for Oracle Secure Backup
     - Output buffer count & size for disk backup can be manually adjusted
       - Details in Support Note 1072545.1 - RMAN Performance Tuning Using Buffer Memory Parameters
       - Note: With Oracle Database 11g Release 2 & ASM, all buffers are automatically sized for optimal performance

Performance Tuning Methodology

- RMAN Backup Data Flow
- Performance Tuning Principles
- Diagnosing Performance Bottlenecks

Diagnosing Performance Bottlenecks

Part – 1

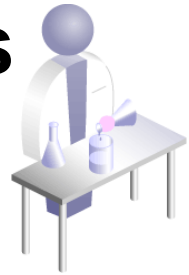

- Query **V$BACKUP_ASYNC_IO**
  - Check **EFFECTIVE_BYTES_PER_SECOND** column (EBPS) for row where **TYPE = 'AGGREGATE'**
  - If EBPS < storage media throughput, run **BACKUP VALIDATE**
    - Case 1: **BACKUP VALIDATE** time ~= actual backup time, then read phase is the likely bottleneck
      - Refer to RMAN multiplexing and buffer usage guidelines
      - Investigate 'slow' performing files: find data file with highest (**LONG_WAITS / IO_COUNT**) ratio
        - If ASM, add disk spindles and/or re-balance disks
        - Move file to new disk or multiplex with another 'slow' file

Diagnosing Performance Bottlenecks

Part – 2

- Case 2: If **`BACKUP VALIDATE`** time << actual backup time, then buffer copy or write to storage media phase is the likely bottleneck
  - Refer to backup compression and encryption guidelines
  - If tape backup, check media management (MML) settings
    - TCP/IP buffer size
    - Media management client/server buffer size
    - Client/socket timeout
    - Media server hardware, connectivity to tape
    - Enable tape compression (but not RMAN compression)

Restore & Recovery Performance

Best Practices

- Minimize archive log application by using incremental backups
- Use block media recovery for isolated block corruptions
- Keep adequate number of archived logs on disk
- Increase RMAN buffer memory usage
- Tune database for I/O, DBWR performance, CPU utilization
- Refer to MAA Media Recovery Best Practices paper
  - Active Data Guard 11g Best Practices (includes best practices for Redo Apply)

Summary

Effective RMAN Backup & Restore Strategy

1. Recovery & business requirements drive the choice
   - Disk?
   - Tape?
   - Data Guard?
2. RMAN tuning: find bottleneck at each phase & remove it
   - Read blocks into input buffers (memory, disk I/O)
   - Copy to output buffers (CPU, compression, encryption)
   - Write to storage media (memory, I/O, media management, HW config)
3. Know your media management product and tape configuration
   - What matters is the end-to-end throughput!

Hardware and Software
Engineered to Work Together


ORACLE®