




ORACLE®



ORACLE®




Oracle Rdb Alternate Backup Techniques

Norman Lastovica
Oracle OpenVMS Development Team

13 October 2009

Agenda

- Background
- Offline backups
- Shadow Sets
- Storage Technologies



Source: Author's estimates based on data from Safeware, The Insurance Agency, Inc., "2000 Safeware Loss Study," 2001; and ONTRACK Data International, Inc., "Understanding Data Loss," 2003.



Background

- IMHO, *backup* performance is generally a non-issue but *restore/recover* performance the actual issue
- In some instances, however, backup performance has a life of its own
 - Database can not be backed up online
 - Exclusive transactions, for example
 - “Replay” repeatable testing



One Offline Fast Backup Technique

- Assumes safe & rapid recovery as ultimate goal
 - Disks are inexpensive – Downtime isn't
- Requires reasonable amount of coordination & care
- Requires off-line Rdb database backups
- Hardware intensive
 - Database remains protected from single point of failure at all times

Disclaimer





Fast Offline Backup Overview

- Utilize host-based shadowing or controller-based “clones” to make “instant” database copies for backup or recovery



Step 1 - Configuration

- Put **all** database files (.RDB, .RDA, .SNP) and **only** database files on shadow set(s)
 - Nothing else on these disks that you can not afford to loose
 - Recovery files (.RDC, .AIJ, .RUJ) separate from database
 - [Remember that live AIJ is most important in database](#)
- All live & snapshot storage areas & root file must all be referred to only by disk volume label (DISK\$foo)
 - Never by physical device name



Step 1 - Configuration

- Organize disks conceptually into 3 “pools”
 1. Two disks for each volume for “online”
 2. One disk for each volume for “backupA”
 3. One disk for each volume for “backupB”
- Logical LUN numbering can be significant win
 - 5xyy, 6xyy, 7xyy, 8xyy for example
- Critical to always keep track of what is on each LUN
 - Changing volume labels can help coordinate this



Step 2 – Prepare for each backup ahead of time

- At some point well prior to backup needing to run
 - Long enough that a “full copy” operation on all volumes will be known to be complete
 - IO load likely higher during shadow copy
- 1. Add third shadow set member to all volumes that hold .RDB, .RDA, .SNP
 - Not needed for .AIJ, .RUJ, .RDC
- 2. Wait for shadow copies/merges to complete



Step 3 – Database close & AIJ backup

- Close database (after an optional AIJ backup) then backup AIJ
1. Online RMU/BACKUP/AFTER_JOURNAL
 2. RMU/CLOSE/WAIT [/ABORT={FORCE|DELPRC}]
 - Database at this point must be closed & consistent
 3. RMU/BACKUP/AFTER_JOURNAL



Step 4 – Split shadow sets & re-open main database

1. Remove third member from each shadow set
 - DISMOUNT /CLUSTER
2. RMU/OPEN/WAIT main database
 - It is now available for access



Step 5 – Privately mount shadow sets

- Mount all shadow set members as process-private
 - Critical step – every member volume that holds database components must be mounted to the process
- **MOUNT /IGNORE=SHADOW /PROCESS**
 - For each member



Step 6 – The Backup

- Do an offline backup from the privately mounted shadow set members
 - Use /NODATABASE_VERIFY
 - Avoid an otherwise expected warning
 - /ALLOCATE=n, /EXTENSION=m & /CRC=CHECKSUM
 - Consider /COMPRESS
- Dismount (or leave mounted to batch job) these members; tomorrow use other set of LUNs
 - There is always a full database online ready to be recovered

3 Member Shadow Set

Backup this
database



DSA700
Database



AIJs



Backups





Recommendations

- Use automated scheduled jobs
 - Too important to leave this up to humans
- Create bomb-proof procedures
- Cycle between shadow member add-in sets
 - By even and odd day of year?
- Copy disk-based backups to tape & store off-site
 - Make sure to encrypt to protect your data



Several Recovery Choices

1. Restore last full backup

- This will take at least as long to run, in most cases, as original backup (sometimes a fair bit longer)
- Hint: Use /SNAPSHOT_ALLOCATION=1

2. Use hot standby database

- Restored & recovered generally in less than 5 minutes of decision to use
- Recovery from certain kinds of failures of DBor hardware

3. Use removed shadow set members

- Restored generally within 5 minutes of decision to use
- Recovery can commence immediately



Recovery using shadow set members

1. Fully dismount original shadow sets
 - Suggest batch job to privately mount one member from each of shadow sets to protect disk content from loss or abuse
2. Re-mount original shadow sets (with original DSA names and volume labels)
 - **From the last removed backup set**
3. Recover database
 - Suggest specifying *.AIJ and /ORDER_AIJ to assist
 - From live & backup AIJs (search-list logical name can help here)
4. Add second member to each shadow set volume



Another layer of protection

- If there are problems with shadow-set method, there is RBF for “traditional recovery”
 - This time can be eliminated by restoring database ahead of time to another set of disks



Downtime is a Bad Time

- Waiting until failure before doing the restore?
- Have restored database from last backup
 - Ready to recover at a moment's notice
 - Verify this database after restore – proves it and the original are in good shape
- Automated backup procedure?
 1. Backup AIJs
 2. Backup database
 3. Restore database
 4. Verify restored database



Simplified Disk Volumes

- Use host-based RAID-0 (ie, software RAID) arrays as to distribute IO load over controllers, disks & adaptors
 - Performance advantages
- Put entire database (.RDA, .RDB, .SNP) on single shadow set
 - Only one member to add/remove/keep track of
 - Downside is only on active HBA or volume



Practice Practice Practice

- Automate with procedures to prevent mistakes during times of stress
- Practice!
 - Work on sequences of events and timings so that you'll know what to do and how long it should be expected to take when it really matters



EVA and XP Business Copy Features

- Can be used in similar fashion as shadow set members
- Use command-line interface to control actions of creating copies
- Still requires closed database to make copy

Business Copy Features

- Virtually instantaneous Snapclone
- Capacity-free Snapshot (Vsnap)
- Immediate use of clone copy, saving significant time and provides a very efficient, low cost local data replication option.





Other Similar Uses

- Repeatable (re)testing with known starting point database state



“Rolling Recoveries”, “Warm Standby”, “Poor Man’s Standby”

- Initial full backup/restore
- Periodically RMAN quiet-point AIJ backup
- Recover “warm” database with AIJ backup
- Remember safe & rapid recovery ultimate goal



RMU /OPTIMIZE /AFTER_JOURNAL

- Can make recovery step(s) faster
- Test with your database & application




Summary

- Not impossible and not difficult
- Requires (briefly) closed database
- Planning, Practice and Procedures are critical




For More Information

search.oracle.com



or

oracle.com/rdb



The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



ORACLE IS THE INFORMATION COMPANY