ORACLE

High Speed Database Recovery

A case study following two recent customer engagements

Agenda

1. Problem Statement
2. Solution
3. Preparation
4. Database Restore
5. Troubleshooting a failed Restore session
6. Conclusion

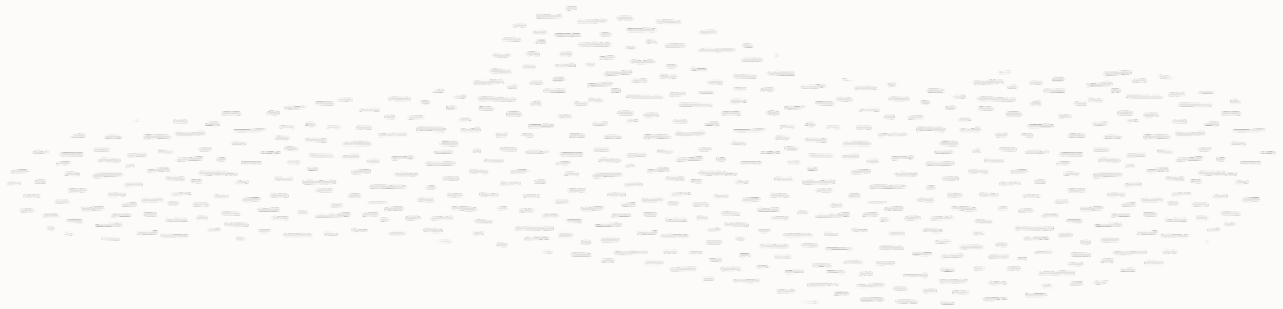

Problem Statement

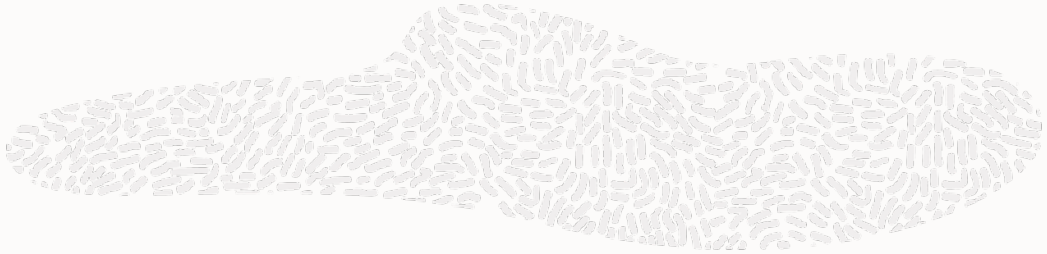

Databases are getting larger

- 10TB Databases are very common
- 100TB+ Databases are common
- 1PB+ Databases are still rare but are becoming part of the conversation

Backup & Recovery Windows continue to shrink

RMAN by default uses only 1 instance of a RAC database during database restore and recovery

- RMAN Duplicate Database command is limited to a single instance

Restore operations are limited by network infrastructure

- Both Ethernet and Storage networks

System Resources are idle during critical operations

- All available cluster wide system resources are not leveraged during the database restoration

Solution : Use all available resources

Documentation, Preparation, and Practice are Essential

Solution: Use all available resources

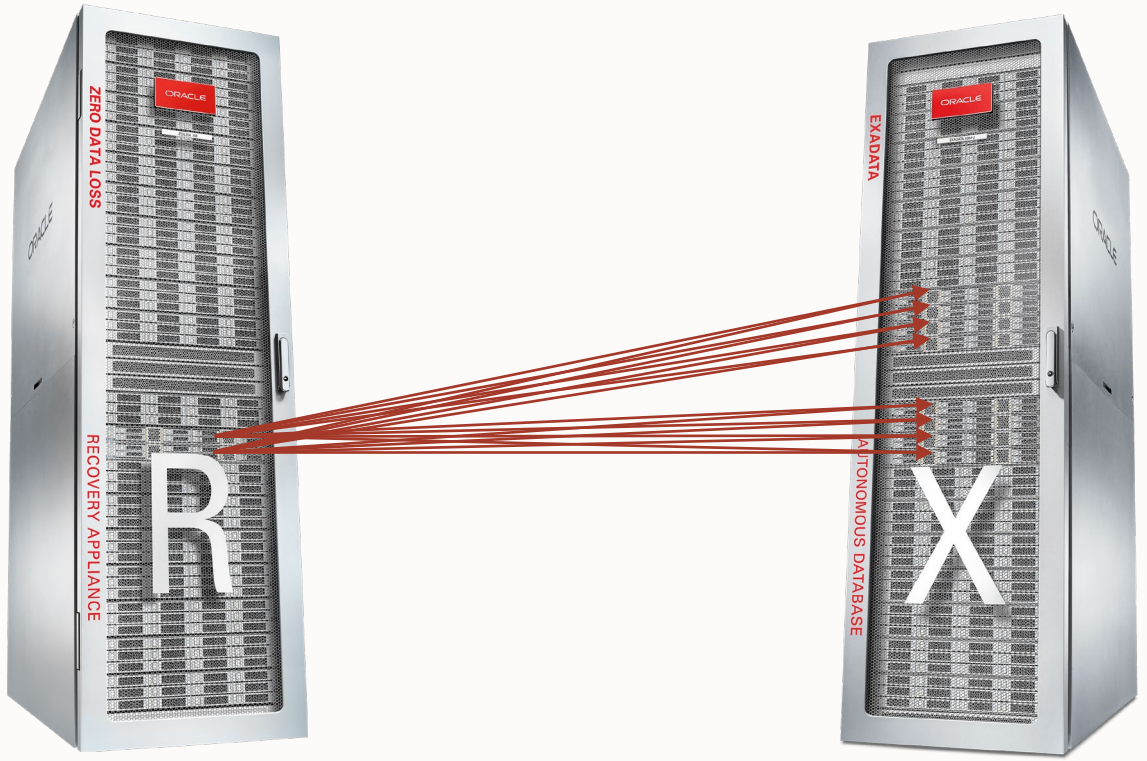

Zero Data Loss Recovery Appliance

Exadata

Solution: Use all available resources

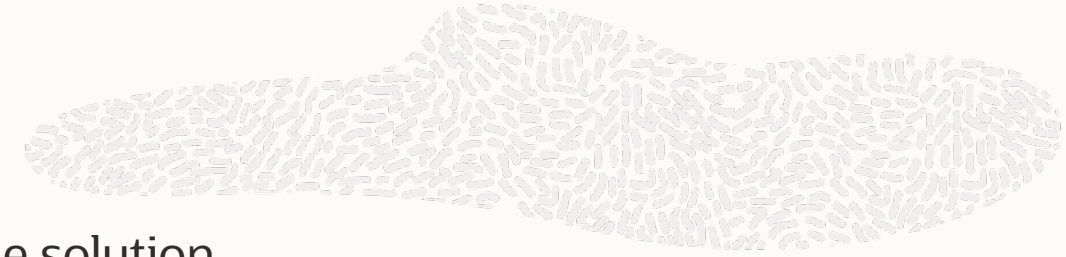

The following data was captured while building the sample solution

Number of Compute Nodes	RMAN Channels per Node	Average Elapsed Restore Time (hr:mi:ss)
1	64	7:55:40
2	32	3:55:27
4	16	2:03:01
8	8	1:16:12

Everything we will discuss here is documented in the MOS Note

- "HOWTO - Improve RMAN restore database performance by using Multi-Instance Database Restore (Doc ID 2710709.1)"

Preparation

Preparation is the Key to Success

Preparation

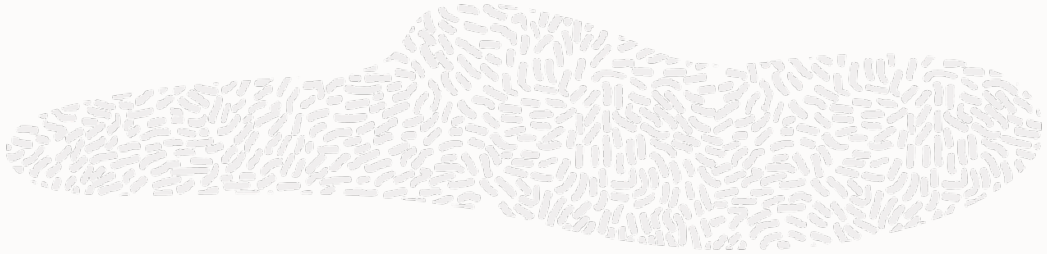

Build a database init.ora file that can be restored when required.

HINT: Review periodically for changes and update accordingly

If the database is Encrypted, know where to restore the encryption wallet from.

HINT: Oracle Key Vault might help eliminate extra copies of encryption keys

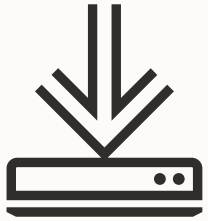

Know how much data should be restored.

HINT: The ZDLRA provides the latest SCN for the current recovery window

GOOD NEWS: If this is a periodic refresh, then the init.ora and encryption setup will probably not change between refreshes

Database Restore

Documentation, Preparation, and Practice are Essential

Database Restore

Restore the init.ora / spfile

Good News: May only need to be performed once. Create an RMAN Command File with the restore command defined

Restore the encryption wallets or configure Oracle Key Vault

Good News: Only needs to be performed once

Restore the Control File using the SCN from prior step

Hint: Create an RMAN Command File with the RMAN restore command defined

Clear online and standby logfiles to avoid noise in alert log

HINT: Create a SQL*Plus script that clears all available log files

Database Restore

Use Secure External Password Store to hide the passwords and retrieve them when needed

- Especially useful when the target database needs to be refreshed periodically
- Add credentials for the target database and for each instance in the target system
- Add credentials for VPC User if ZDLRA is to be used

Good News: Only needs to be performed once or when the passwords changed

Hint: Oracle Key Vault can also store passwords and secrets

Database Restore

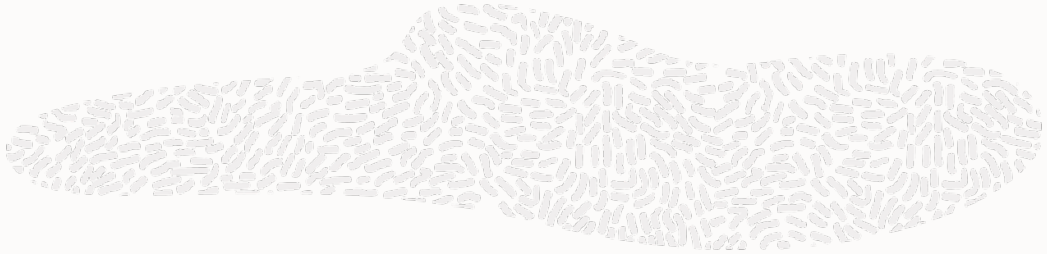

Finally, Almost – a database restore command

- Create an RMAN Command File for the core database restore and recovery commands

```
set echo on
run {
allocate channel SBT_01 device type sbt FORMAT '%d_%U' PARMS "SBT_LIBRARY=<ORACLE_HOME>/lib/libra.so,
SBT_PARMS=(RA_WALLET='location=file:<wallet_location> credential_alias=<ra_credential_alias>')"
connect '/@<destination_scan>/dwdbstby/dwdbs1';
allocate channel SBT_02 device type sbt FORMAT '%d_%U' PARMS "SBT_LIBRARY=<ORACLE_HOME>/lib/libra.so,
SBT_PARMS=(RA_WALLET='location=file:<wallet_location> credential_alias=<ra_credential_alias>')"
connect '/@<destination_scan>/dwdbstby/dwdbs2';
allocate channel SBT_03 device type sbt FORMAT '%d_%U' PARMS "SBT_LIBRARY=<ORACLE_HOME>/lib/libra.so,
SBT_PARMS=(RA_WALLET='location=file:<wallet_location> credential_alias=<ra_credential_alias>')"
connect '/@<destination_scan>/dwdbstby/dwdbs3';
allocate channel SBT_04 device type sbt FORMAT '%d_%U' PARMS "SBT_LIBRARY=<ORACLE_HOME>/lib/libra.so,
SBT_PARMS=(RA_WALLET='location=file:<wallet_location> credential_alias=<ra_credential_alias>')"
connect '/@<destination_scan>/dwdbstby/dwdbs4';
...
...
```

Database Restore

Finally, Almost – a database restore command

- Create an RMAN Command File for the core database restore and recovery commands (continued)

```
...
...
set autolocate off;
set newname for database to new;
set until scn 12593370;
restore database;
recover database;
}
```

set autolocate off [default is on]
Should only be set to off when you are certain that all instances of the cluster involved in the restore can access every backup piece, as with ZDLRA or Cloud backup locations.

set until scn 12593370;
Retrieved from the query performed against the ZDLRA
SQL> select final_change# from rc_database where name = '<dbname>';

Database Restore

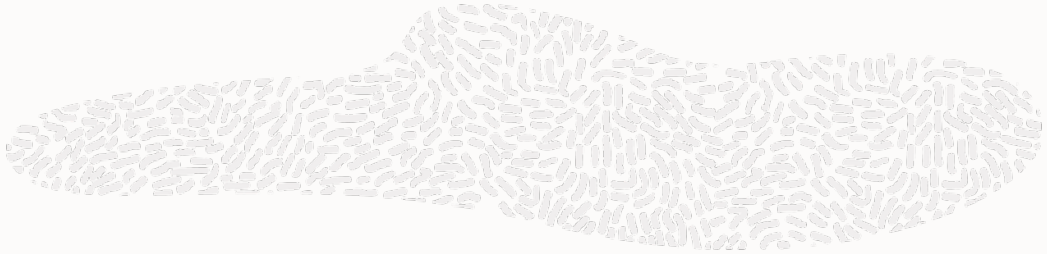

Finally

- Run RMAN command using all the prior steps

  $ rman target /@<destination_scan>/dwdbstby catalog /@<ra_credential_alias> cmdfile=restore_recover_database.rman log=restore_recover_database.log append

  - NOTE: No passwords are visible from the command line or from the RMAN Command File script
    - Therefore, no passwords are present in the BASH history file
    - Using the SCAN address "target /@<destination_scan>/dwdbstby" will leverage all instances in the cluster in one of two ways:
      - If "connect '/@<destination_scan>/dwdbstby/dwdbs2';" is specified in the channel allocations, those channels will be explicitly assigned to that instance
      - If "connect '/@<destination_scan>/dwdbstby/dwdbs2';" is not specified, RMAN will automatically allocate channels across available instances, but the channels will not be equally balanced across all the available instances
      - If "target /" is used and "connect '/@<destination_scan>/dwdbstby/dwdbs2';" is not specified, then only the local instance will be used
    - Log file contains all the output from the session
  - HINT: Ensure the SSH session or VPN session will not timeout – use linux "nohup" if necessary
  - HINT: Set SHELL variable NLS_DATE_FORMAT='DD-Mon-YYYY HH24:MI:SS' to get accurate times

Troubleshooting

Failures Happen, be Prepared and Document Past Failures

Troubleshooting a failed Restore session

MOS Note 2710709.1, has a collection of troubleshooting steps to handle failed restore operations

- Determine the cause of the failure – process death, network errors, disk errors
  - Don't just re-run
  - Open a support ticket immediately
- Detect files that have been successfully restored, and use them
  - And avoid unnecessary duplicate restores
- Detect files that have not been restored completely, and remove them
  - Free-up space in the ASM Disk Group or volume where the files are being restored to

Summary

High Speed Database Recovery Possible with Appropriate Documentation, Preparation, and Practice

Conclusion

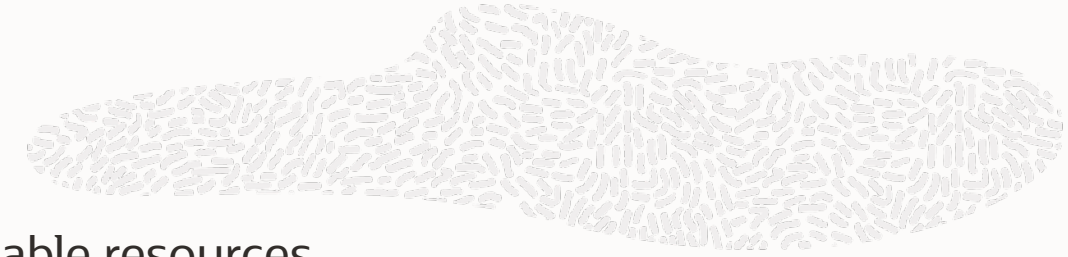

High Speed Database Recovery is possible … Use all available resources

- But it does not come for free
  - Documentation
    - Have a Disaster Recovery Document that all DBAs and support personnel have access to
  - Preparation
    - If you don't prepare, it will likely compound the disaster
  - Practice
    - Don't "Practice" at 2:00am when the corporate database has crashed
    - If you don't "Practice", your probability of success in a crisis is roughly 0
  - Failures will occur
    - But if you have practiced, you may already know the solution
    - Open an SR
      - Include the Database Version, RMAN Command File script and RMAN Log File output
      - Refer to MOS Note "SRDC - Required Diagnostic Data Collection for RMAN Issues (Doc ID 1671431.1)"
- Refer to "HOWTO - Improve RMAN restore database performance by using Multi-Instance Database Restore (Doc ID 2710709.1)"

Additional Resources

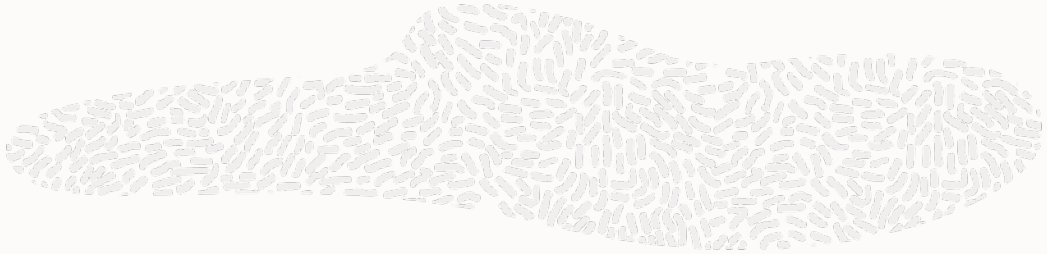

Oracle Support Documents

- HOWTO - Improve RMAN restore database performance by using Multi-Instance Database Restore (Doc ID 2710709.1)
- SRDC - Required Diagnostic Data Collection for RMAN Issues (Doc ID 1671431.1)

Product Information

- Oracle Exadata :
  - https://www.oracle.com/engineered-systems/exadata/
- Oracle Zero Data Loss Recovery Appliance Family :
  - https://www.oracle.com/engineered-systems/zero-data-loss-recovery-appliance/
- Oracle Maximum Availability Architecture (MAA)
  - https://www.oracle.com/database/technologies/high-availability/maa.html

Social

Comment, Like, Share!

Engineered Systems & Storage	SPARC & Solaris
facebook.com/OracleITInfrastructure	facebook.com/OracleITInfrastructure
twitter.com/Infrastructure	twitter.com/Infrastructure
#ITInfrastructure #CloudReady #Exadata # ZDLRA #OracleZFS #ODA	#oraclesparc #sparc #oraclesolaris #solaris
Oracle.com/engineered-systems/ Oracle.com/Exadata Oracle.com/ODA Oracle.com/ZDLRA Oracle.com/storage	Oracle.com/sparc Oracle.com/solaris


ORACLE

Our mission is to help people see data in new ways, discover insights, unlock endless possibilities.