Oracle Maximum
Availability Architecture

Oracle Cloud Infrastructure Exadata Backup & Restore Best Practices using Cloud Object Storage

ORACLE WHITE PAPER | MAY 2019

ORACLE

Contents

Introduction 2

Cloud Backup and Restore Life Cycle Best Practices 3

RMAN Backup and Restore Configuration Best Practices 4

Observed and Expected Backup and Restore Performance 6

- RMAN Full backup 7
- RMAN Incremental backup 8
- RMAN Restore from Full Backup 9

Backup and Restore Maintenance Activities 9

Conclusion 12

Appendix A 13

- RMAN Full backup 13
- RMAN Incremental backup 14
- RMAN Restore from Full Backup 15

Appendix B 16

Appendix C 18

Appendix D 19

Introduction

Oracle Database Backup Cloud Service is a secure, scalable, on-demand storage solution for backing up Oracle on-premises databases or Cloud databases residing in Cloud@Customer or Oracle Cloud Infrastructure (OCI).

Oracle Exadata is Oracle's best Maximum Availability Architecture (MAA) database platform in the Oracle Cloud Infrastructure (OCI).

This paper provides MAA recommended backup and restore best practices for backing up an Exadata Database systems to the OCI Database Backup Cloud Service. This paper will cover:

1. Cloud Backup and Restore Life Cycle best practices
2. RMAN Backup and Restore Configuration best practices
3. Observed and Expected Backup and Restore performance
4. Examples to achieve higher backup and restore performance

Cloud Backup and Restore Life Cycle Best Practices

1. Understand the Backup, Restore and Data Retention Service Level Agreements

The following questions should be addressed before proceeding:

a. What is the expected backup window?

b. What is the expected and maximum restore and recovery time or Recovery Time Objective (RTO) for disaster recovery?

c. What are the data retention requirements?

d. What is the data loss tolerance or Recovery Point Objective (RPO)?

Most Exadata customers use the pre-configured Exadata setup with 80% DATA and 20% RECO disk group configuration. Most Exadata customers also backup to Oracle Cloud Infrastructure object storage since it will provide protection from database disaster such as unrecoverable database, cluster, Exadata rack, or Availability Domain failures. If RTO requirements cannot be met consistently by restoring a backup, MAA recommends a standby database or Golden Gate replica in another Oracle Cloud Infrastructure database resource. This paper provides guidance and MAA best practices on achieving optimal backup and restore rates with insight on varying CPU overhead. The examples provided help estimate backup and restore time when using Oracle Cloud Infrastructure object storage with the Database Backup Cloud Service.

2. Verify the latest Cloud DBaaS tool version is installed

Cloud DBaaS tool RPM release 18.2.3.1.0_190328.0930 or higher incorporates the validated MAA backup/restore best practices. Refer to Updating the Cloud Tooling on Exadata Cloud Service (https://docs.cloud.oracle.com/iaas/Content/Database/Tasks/exatooling.htm#UpdatingToolingonanExadataDBSystem) to update to the latest Cloud DBaaS tool RPM.

3. Determine Data Retention Requirements which will dictate how long backups should be kept in Database Backup Cloud Service object storage

   This is described as bkup_oss_recovery_window in “Backing Up an Exadata Database” documentation:

   https://docs.cloud.oracle.com/iaas/Content/Database/Tasks/exabackingup.htm

The default for bkup_oss_recovery_window is 30 days and can be adapted as needed up to a maximum of 90 days. Also, the Oracle Cloud Infrastructure will automatically set the CONTROL_FILE_RECORD_KEEP_TIME. Backups will automatically written to all Availability Domains.

4. Follow the Backup Life Cycle Best Practices for Oracle Cloud Infrastructure object storage described here: https://docs.cloud.oracle.com/iaas/Content/Database/Tasks/exabackingup.htm

   MAA recommends:

   - Weekly Full Level 0 (default Sunday)

- Daily Incremental (default approach)
- Note that all backups in the Oracle Cloud Infrastructure object storage are encrypted via TDE key (automatically) or password in case of manual backups.

The recommended day for the default level 0 backup and backup start time can be changed (bkup_daily_time=*hh:mm)*. An on-demand backup can also be invoked at any time. Refer to https://docs.cloud.oracle.com/iaas/Content/Database/Tasks/exabackingup.htm for more information.

5. Validate Backup and Restore Periodically

There are automatic and manual steps that MAA recommends on a periodic bases:

- Daily RMAN crosscheck. This is performed automatically by any Cloud database backup that is configured by default, invoked by the cloud console or invoked by the Cloud DBaaS tooling.
- Monthly *restore validate* operation with *check logical* option. This must be performed manually. Pick a time window that can tolerate an additional 8% CPU utilization overhead. This operation does not interfere with running backups. For an example, refer to the *Backup and Restore Maintenance Activities* section of this paper.
- Quarterly full restore and recovery evaluation. This must be performed manually. For an example, refer to the *RMAN Restore from Full Backup* section of this paper.

RMAN Backup and Restore Configuration Best Practices

This section describes the key recommended backup and restore settings and the impact on the overall backup and restore performance. The settings below are already the default in the DBaaS RPM 18.2.3.1.0_190328.0930 or higher release. If the default settings need to change, refer to *bkup_api get config* or *set config command* in Appendix D.

1. Use Default RMAN Section Size=64GB:

Oracle does not recommend changing this setting. The purpose of multi-section backups is to enable RMAN channels to back up a single large file in parallel. RMAN divides the work among multiple channels, with each channel backing up one file section in a file. Backing up a large data file in separate sections can significantly improve the performance of backups and the setting will be ignored or auto-tuned for smaller datafiles. Note: Oracle 11g databases can restore an initial level 0 multi section but not the incremental level 1 backups.

2. Use default RMAN parallelism and tune if required

The default total RMAN Channels count is 8 for Exadata Quarter Rack, 16 for Exadata Half Rack and 32 for Exadata Full Rack. RMAN channels are spread across the RAC database instances. The default provides decent backup performance with very minimal database performance impact (<5% CPU utilization); however if this is the only database backup occurring for specific window of time, changing the RMAN channel count to 16 will provide more backup performance with tolerable database (<5% CPU utilization) impact. For higher throughput with additional CPU utilization, continue to increase the RMAN channels by a factor of 2. The maximum total RMAN channels per Exadata system is 64; the maximum RMAN channels for one Exadata node is 32. Refer to the performance examples provided later in this paper for guidance. If the full number of CPUs available to Exadata database server are not being used, reduce the total RMAN channel count to limit backup overhead. Use the guidelines in the table below for channel recommendations:

Total CPU Count per node	Recommended Total RMAN channel for Exadata Quarter Rack
46 or higher	8 channels (Default) or more
20 > #cores < 46	4 to 8 channels
Less than 20	2 to 4 channels

3. Default RMAN compression is LOW

By compressing backups prior to sending to Oracle Cloud Infrastructure object storage, the overall backup storage used, elapsed backup time and network utilization will be reduced. The potential trade-off is additional CPU utilization during backup and restore operations. MAA findings show that an RMAN compression of LOW provides respectable backup throughput with reduced Oracle Cloud Infrastructure object storage capacity, while still keeping Exadata DB node CPU overhead under 5%. To reduce CPU utilization during backup and restore operations, disable RMAN compression and issue RMAN commands outside the cloud APIs. Cloud APIs allow changes to different levels of RMAN compression but currently setting to NONE is not available via cloud tooling or *bkupapi set config command*. Be aware that setting compression to NONE will increase backup storage utilization in the Oracle Cloud Infrastructure object storage. Refer to Appendix A for uncompressed backup performance examples when not using the Cloud API.

If tablespaces are configured with Exadata Hybrid Columnar Compression (HCC) compression, backups with LOW compression can incur additional CPU overhead. Disabling RMAN compression during database backups (not archivelog backups) for those target tablespaces, will reduce the CPU overhead and still obtain respective backup and restore throughput with reduced Oracle Cloud Infrastructure object storage utilization. Refer to Appendix C for performance tests and RMAN command examples.

4. Keep the control file retention the same as the backup retention

The control file retention time is automatically adjusted to the same value as the object storage retention window or `bkup_oss_recovery_window` plus 8 days. The maximum `bkup_oss_recovery_window` is 90 days.

5. Default archive backup frequency is 1 hour to minimize RPO potential

The `bkup_archlog_frequency` setting ensures that the RMAN backup cron job is initiated to back up any archives that have not been backed up to the Oracle Cloud Infrastructure object storage. This limits data loss or RPO exposure to a little over an hour. Starting with Cloud Release 18.2.3.1, there is an option to change this value and reduce it to a minimum of 15 minutes, which will further reduce RPO exposure. Note that this does not cover the changes or data found in the Online Redo Logs that are not archived. MAA best practices also dictate a low log switch frequency (once every 15 minutes) to optimize database performance. To achieve zero or sub-second level RPO or a more guaranteed RPO < 30 minutes, MAA recommends a configuration with Data Guard and separate Exadata in Oracle Cloud Infrastructure to host the standby database. Refer to Using Oracle Data Guard in Exadata Cloud Service (https://docs.oracle.com/en/cloud/paas/exadata-cloud/csexa/use-data-guard-this-service.html#GUID-FB63CBB2-74E1-4FD6-AFFA-4F902008223C).

Observed and Expected Backup and Restore Performance

These performance results were observed with the following configuration:

- Cloud 18.2.5.1 or higher
- Exadata Quarter Rack or Higher
- Exadata version 18.1.6 or higher
- Database version 12.2 (12.2.0.1.180116) or higher

The Backup and Restore results shown below are based on testing performed against a 12.8TB container database with six pluggable databases (four pluggable databases did not have any compression; one pluggable database was compressed for OLTP and one pluggable database was compressed for HCC Query High) running on an Exadata X7-2 quarter rack (92 CPU cores and 742 GB of RAM per database node).

A Swingbench workload with ~2000 - ~2500 TPS and 5% - 8% daily change rate was used against the database. *The backups were executed while the workload is running.*

Note that backup and restore throughput may vary depending on the change rate and data distribution but these examples will provide some potential backup/restore performance expectations.

RMAN Full backup

Figure 1 below shows the results of the RMAN database full backup tests. The tests were performed using various RMAN channels and the default LOW compression (See Appendix A). By default, the RMAN full backups occurs once a week. To obtain a lower CPU overhead the number of channels can be reduced, but this will extended the backup window.

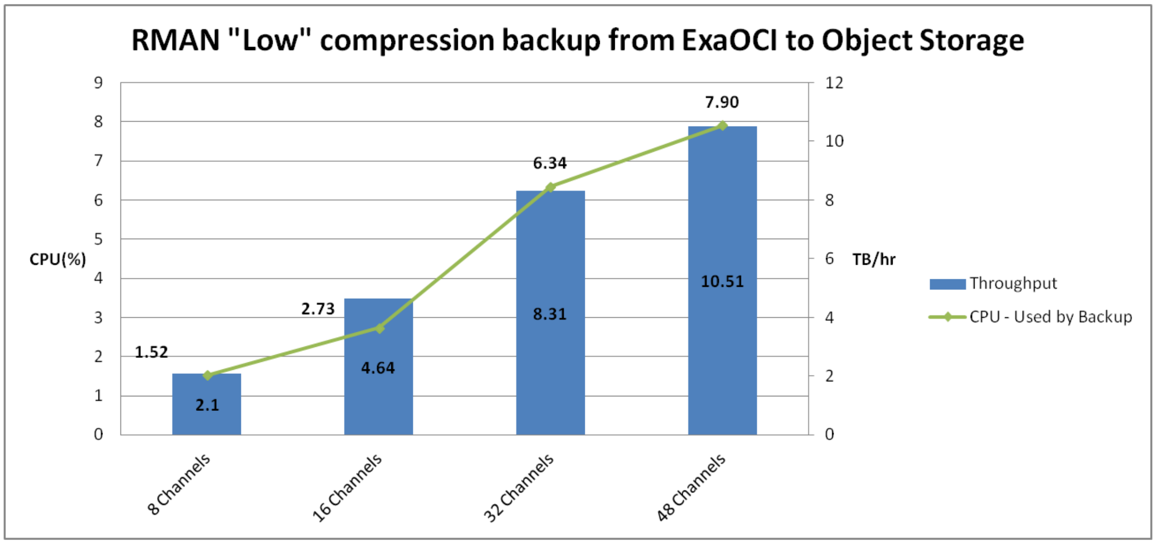

Figure 1: RMAN full database backup

RMAN full backup performance observations with various parallelism settings:

With the default RMAN parallelism setting of 8 total channels (4 per node), 2.1 TB/hour with 1.52% CPU overhead from the backup operation.

With RMAN parallelism of 16 total channels and LOW compression, a throughput of 4.64 TB/hour with 2.73% CPU overhead from the backup operation → recommended for higher backup throughput.

By increasing RMAN channels to 48 or more total channels, a throughput of over 10.51 TB/hour with < 8% CPU overhead.

Please note the above total channel recommendations is targeted for full Exadata Quarter Rack utilizing all 92 CPU cores for Exadata Database Server. If the total enabled CPU cores is significantly less, it will be necessary to reduce the total number of RMAN channels.

Conclusion: The backup throughput will increase linearly when adding more RMAN channels. However, by increasing the number of RMAN channels, more CPU overhead is observed. It is recommended to keep the total number of RMAN channels between 8 and 96.

RMAN Incremental backup

The goal of an incremental backup is to back up only those data blocks that have changed since a previous backup. Incremental backup throughput is referred to as “effective throughput” or the time taken to back up the incremental changes needed for the database recoverability divided by the database total size. By default, Incremental backups occur daily except on the designated full backup day.

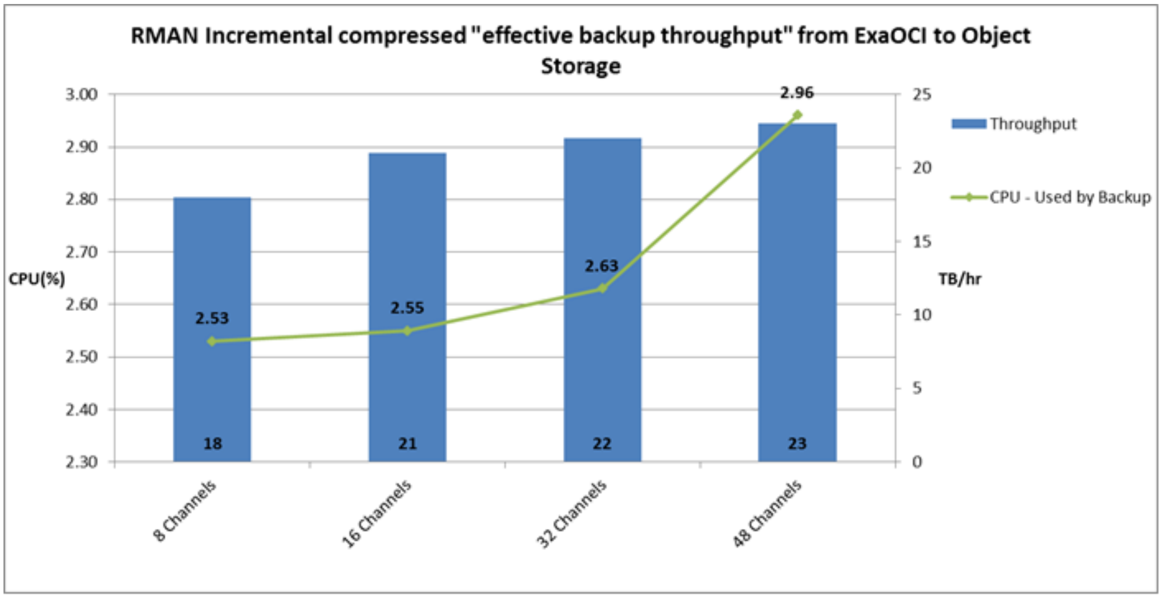

Figure 2: RMAN incremental database backup

RMAN incremental backup performance observations with various parallelism settings:

With RMAN parallelism of 16 total channels and LOW compression, an effective throughput of 21 TB/hour with 2.55% CPU overhead from the backup operation.

Effective throughput increases to 23 TB/hour but CPU overhead increases to 2.96% as the total number of channels increases to 48.

Please note the above total channel recommendations is targeted for full Exadata Quarter Rack utilizing all 92 CPU cores for Exadata Database Server.

Conclusion: The default backup strategy and configuration incur very little overhead (<2%) while resulting in very high daily incremental backup rates (18 TB/hour) and tolerable full backup rates (2.1 TB/hour).

RMAN Restore from Full Backup

During Restore and Recovery operations, the number of RMAN channels can be increased in order to reduce the Database RTO. The goal is to maximize the throughput with the adequate number of channels.

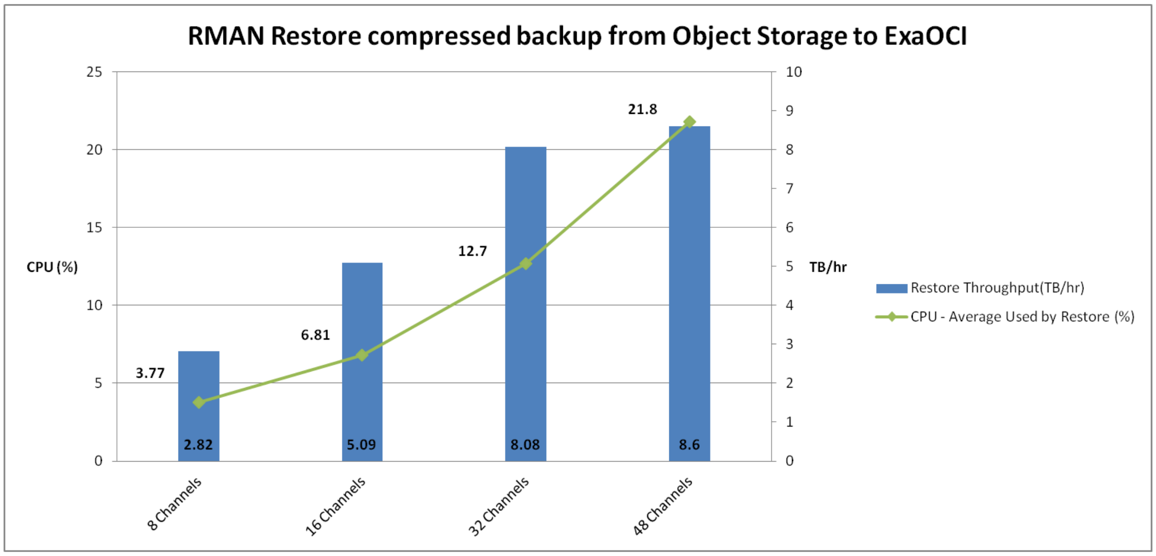

Figure 3 RMAN full database restore

RMAN restore performance observations:

With RMAN parallelism of 48 total channels, a restore throughput of 8.6 TB/hour with 21.8% CPU utilization → recommended.

The default RMAN parallelism used by Cloud Tooling for restore is too low.   At this time, MAA recommends increasing the total number of channels during database restore operation.

Conclusion: The goal during restore operations is to maximize the throughput to reduce the RTO. It is recommended to use 48 channels which requires an explicit RMAN restore operation outside Cloud Tooling.

Backup and Restore Maintenance Activities

The recommended RMAN maintenance activities are backup crosscheck and validation. The Cloud database backup API automatically executes a crosscheck and delete obsolete with backup. RMAN restore validate must be manually executed.

This important automated crosschecking of Cloud backups occurs with the default daily backup and introduces very little overhead. If a backup set, or a piece, is missing it should be deleted. Crosschecking only marks the missing backup sets/pieces as expired and does not delete or remove anything from object storage. Backup sets/pieces marked as expired will not count toward the retention policy of the “delete obsolete” command.

The recommended restore validate operation should be manually performed at least once a month.

RMAN restore validate reads the backup sets and checks for corruption. The RMAN restore validate command acts as a normal restore operation but without the overhead of writing any data to database storage. The data is streamed from the backup in object storage to the database for validation purposes and gets discarded after the validation. The Cloud Database backup API provides the RMAN validate option using the reval_start argument.

Refer to Appendix B for performance examples when not using Cloud API.

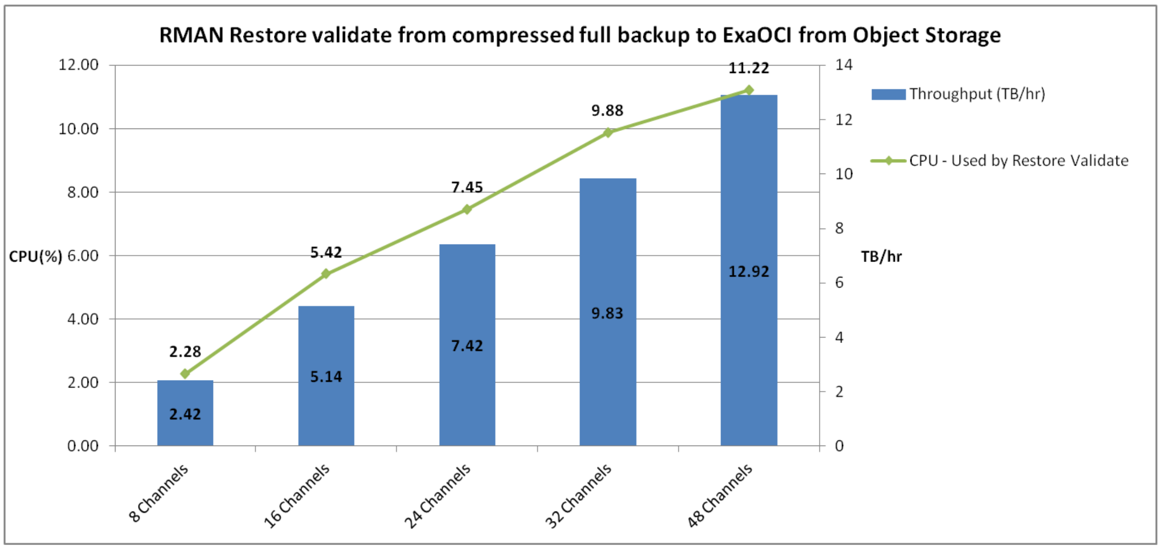

RMAN restore validate from compressed full backup performance observations with various parallelism settings:

With default RMAN parallelism of 8 total channels, a throughput of 2.42 TB/hour with 2.28% CPU overhead from the restore validate operation.

If additional CPU headroom is available without impacting application performance, RMAN parallelism can be increased. Performance examples are below:

RMAN parallelism of 24 total channels, a throughput of 7.42 TB/hour with < 8% CPU overhead.

Increasing RMAN total channels to 48, a throughput of 12.92 TB/hour with < 12% CPU overhead.

The RMAN restore validate command does a block level check of the backups and ensures that a restore could be performed, by confirming that all database files exist and are free of physical and logical corruptions. If you have large data retention requirement, it is recommended to validate using a date restriction as specified below.

```
restore database validate until time "to_date('2019-03-07 22:58:37','YYYY-MM-DD
HH24:MI:SS')";
```

or

```
restore database validate until time "SYSDATE - 7";
```

Finally, the recommended complete restore and recovery tests should be manually performed once a quarter.

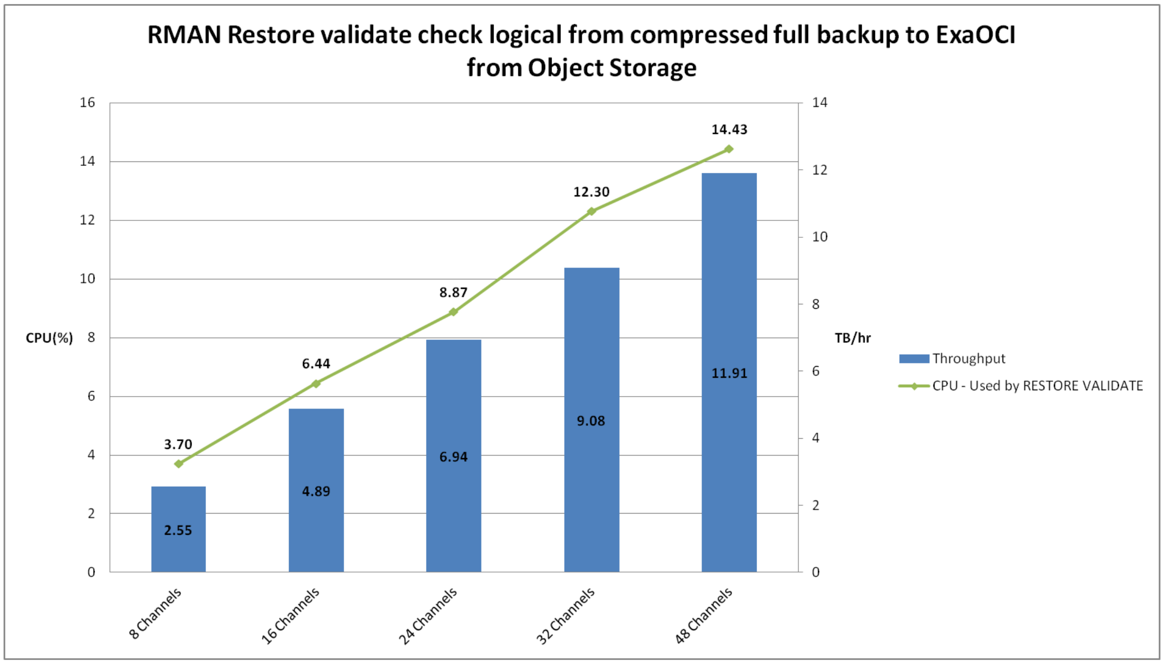

RMAN compressed restore validate check logical performance observations with various parallelism settings:

With default of 8 total channels, a throughput of 2.55 TB/hour with 3.70% CPU overhead from the restore validate operation.

With RMAN parallelism of 16 total channels, a throughput of 4.89 TB/hour with 6.44% CPU overhead.

By increasing RMAN total channels to 64, a throughput of 15.13 TB/hour with 19.95 % CPU overhead.

Refer to Appendix B for examples.

Conclusion: Backup validation is important, and it is recommended to be done monthly. This activity should be performed at a window that can tolerate the CPU overhead. To reduce the time of the maintenance activity, the number of channels can be increased but the CPU overhead will also increase.

Conclusion

Oracle Cloud Database Backup Service is an effective and low-cost solution to protect Oracle databases. By leveraging the MAA configuration and operational practices, the restore and recovery operations from Oracle Cloud object store will be more successful and predictable.

Appendix A

The tests in this appendix were executed without RMAN compression. The DBaaS API or cloud tooling cannot be used for uncompressed backup/restore operations. Issue the RMAN commands directly. Refer to Appendix D.

RMAN Full backup

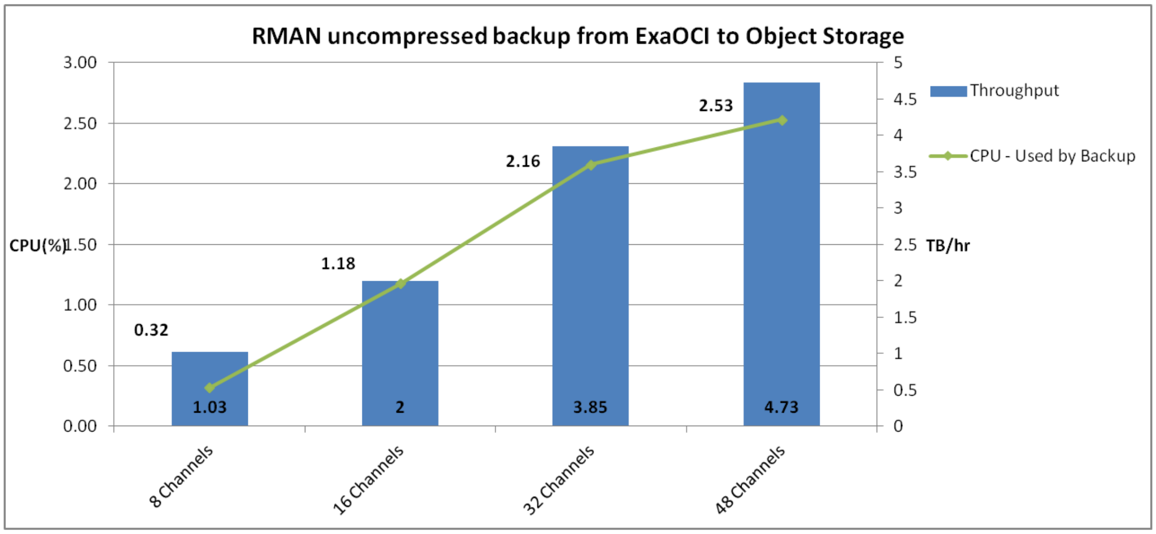

RMAN uncompressed full backup performance observations with various parallelism settings:

With RMAN parallelism of 16 total channels, a throughput of 2 TB/hour with 1.18% CPU overhead.

By increasing RMAN channels up to 48 total channels, a maximum of 4.73 TB/hour with < 3% CPU overhead.

Sample RMAN level 0 backup script:

```
RUN
 {
 CONFIGURE DEVICE TYPE 'SBT_TAPE' BACKUP TYPE TO BACKUPSET PARALLELISM 32;
 BACKUP DEVICE TYPE SBT INCREMENTAL LEVEL 0 TAG ‘MY-TAG’ AS BACKUPSET SECTION SIZE 64G DATABASE
PLUS ARCHIVELOG NOT BACKED UP FORMAT '%d_%U' DELETE INPUT;
 }
```

RMAN Incremental backup

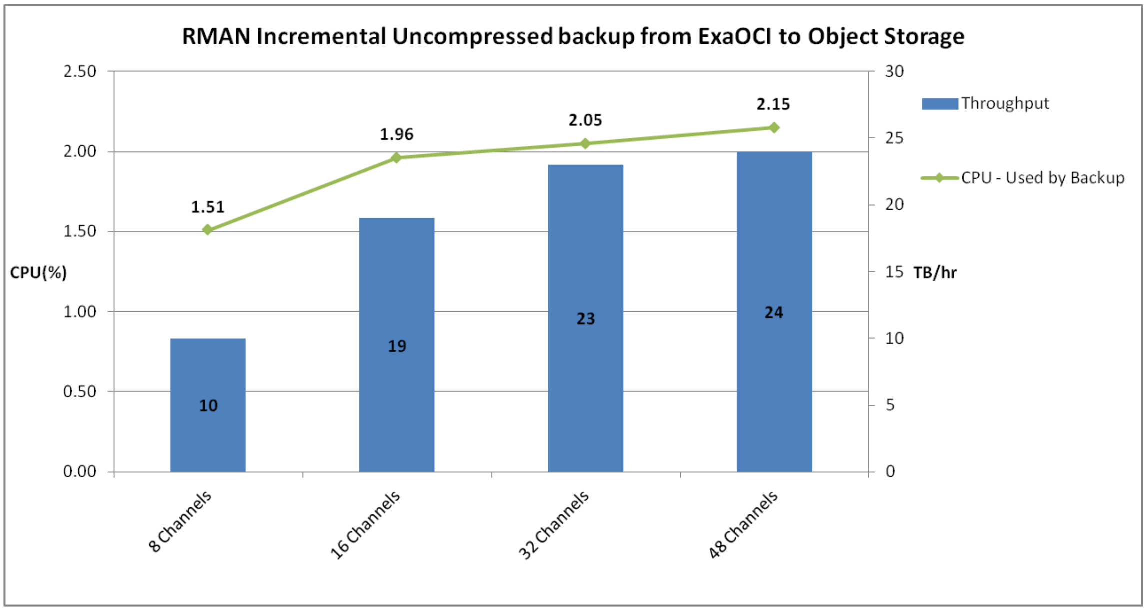

RMAN uncompressed incremental backup performance observations with various parallelism settings:

With RMAN parallelism default of 16 total channels, an effective throughput of 19 TB/hour with < 2% CPU overhead.

With RMAN parallelism of 32 total channels, an effective throughput of 23 TB/hour with < 3% CPU overhead → recommended.

By increasing RMAN channels up to 48 total channels, a maximum effective throughput of 24 TB/hour with < 3% CPU overhead.

Sample RMAN level 1 backup script:

```
RUN
 {
 CONFIGURE DEVICE TYPE 'SBT_TAPE' BACKUP TYPE TO BACKUPSET PARALLELISM 32;
 BACKUP DEVICE TYPE SBT INCREMENTAL LEVEL 1 TAG ‘MY-TAG’ AS BACKUPSET SECTION SIZE 64G DATABASE
PLUS ARCHIVELOG NOT BACKED UP FORMAT '%d_%U' DELETE INPUT;
 }
```

RMAN Restore from Full Backup

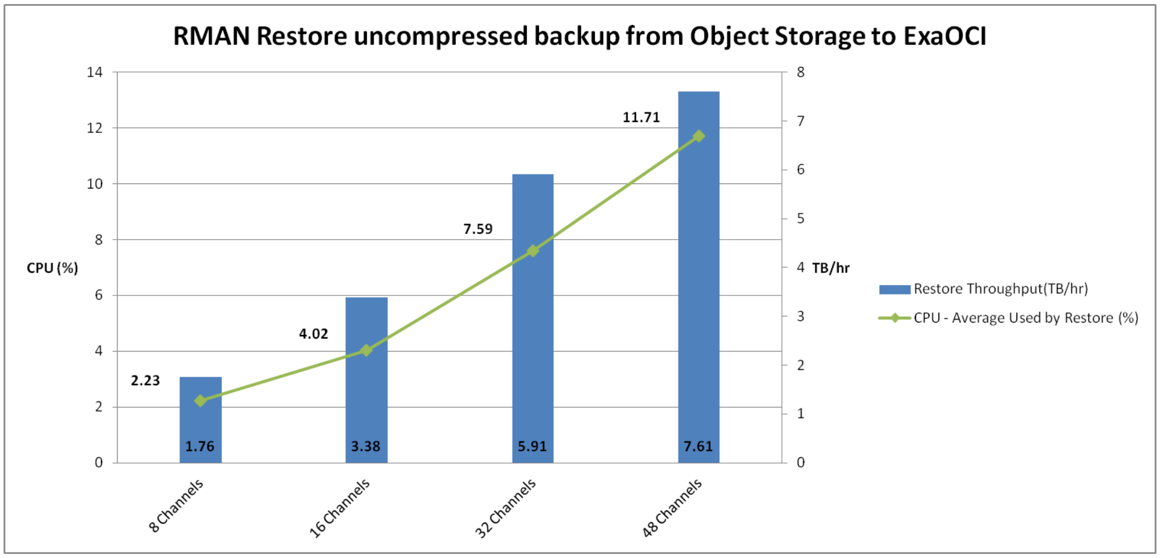

RMAN uncompressed full restore performance observations with various parallelism settings:

With RMAN parallelism of 48 total channels, a restore throughput of 7.61TB/hour with 11.71% CPU utilization.

Sample RMAN restore script:

```
RUN
 {
 CONFIGURE DEVICE TYPE 'SBT_TAPE' BACKUP TYPE TO BACKUPSET PARALLELISM 32;
 RESTORE DATABASE;
 RECOVER DATABASE;
 }
```

Appendix B

The tests in this appendix were executed from a backup without RMAN compression.

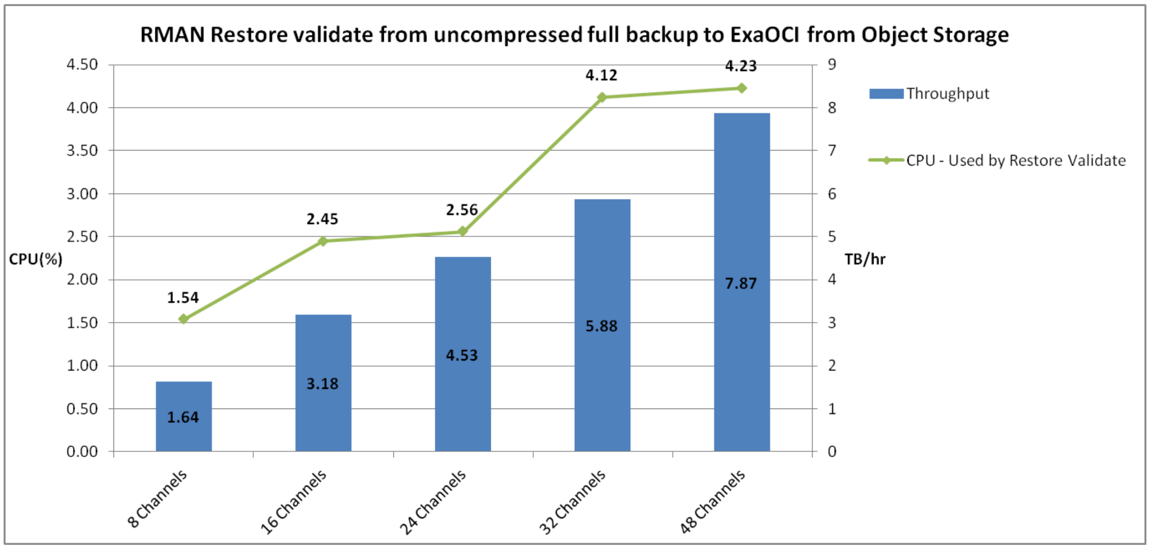

RMAN uncompressed restore validate performance observations with various parallelism settings:

With RMAN parallelism of 32 total channels, a throughput of 5.88 TB/hour with < 5% CPU overhead.

By increasing RMAN total channels to 48, a throughput of 7.87 TB/hour with < 8% CPU overhead.

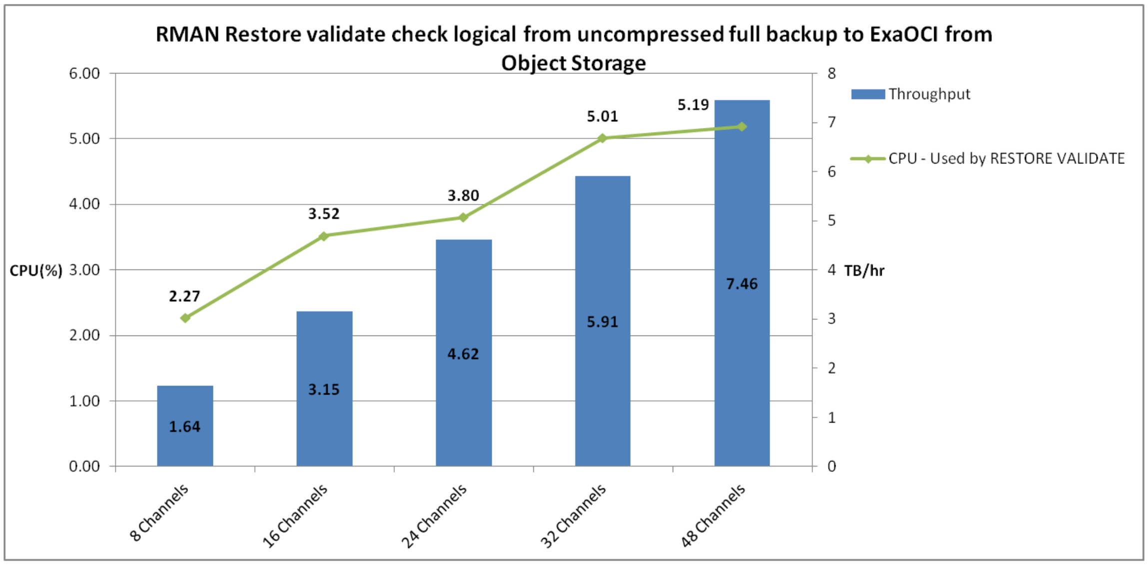

RMAN uncompressed restore validate check logical performance observations with various parallelism settings:

With RMAN parallelism of 24 total channels, a throughput of 4.62 TB/hour with < 4% CPU overhead.

By increasing RMAN total channels to 48, a throughput of 7.46 TB/hour with 5.19 % CPU overhead.

Sample restore validate using the backup API or RMAN script:

 # /var/opt/oracle/bkup_api/bkup_api reval_start --dbname=mydb

For a complete list of options using the cloud backup API, please refer to the documentation:
https://docs.oracle.com/en/cloud/paas/exadata-cloud/csexa/back-and-recover.html

Sample RMAN restore validate code:

```
RUN
 {
 CONFIGURE DEVICE TYPE 'SBT_TAPE' BACKUP TYPE TO BACKUPSET PARALLELISM 48;
 RESTORE DATABASE VALIDATE FROM TAG 'MY-TAG-COMP';
  }
```

Sample RMAN restore validate code – with check logical:

```
RUN
 {
 CONFIGURE DEVICE TYPE 'SBT_TAPE' BACKUP TYPE TO BACKUPSET PARALLELISM 48;
 RESTORE DATABASE VALIDATE CHECK LOGICAL FROM TAG 'MY-TAG-COMP';
  }
```

Appendix C

Level 0 backup tests were performed against pluggable databases: non-compressed data and OLTP compressed data. The table below shows the different metrics observed.

Database	Size (GB)	Database Compression	Total RMAN Channels	RMAN Compression	Duration	Throughput (TB/hr)	CPU (%) - Added by Backup
PDB1	2370.09	None	8	None	9291	0.9	1
PDB1	2370.09	None	16	None	5561	1.5	2
PDB1	2370.12	None	32	None	2202	3.78	6
PDB1	2370.1	None	48	None	2225	3.74	7
PDB1	2370.15	None	8	LOW	5097	1.63	6
PDB1	2370.19	None	16	LOW	1877	4.44	16
PDB1	2370.19	None	32	LOW	1275	6.54	27
PDB1	2370.24	None	48	LOW	1335	6.24	26
PDB4	2461.03	OLTP	8	None	9119	0.95	2
PDB4	2461.41	OLTP	16	None	5448	1.59	3
PDB4	2461.57	OLTP	32	None	2169	3.99	7
PDB4	2460.99	OLTP	48	None	2255	3.84	7
PDB4	2461.08	OLTP	8	LOW	5917	1.46	6
PDB4	2461.39	OLTP	16	LOW	1879	4.61	15
PDB4	2461.47	OLTP	32	LOW	1333	6.49	25
PDB4	2461.54	OLTP	48	LOW	1379	6.28	25

Appendix D

To change the default number of RMAN channels used by the backup API, set the option bkup_channels_node:

Example: `# /var/opt/oracle/ocde/assistants/bkup/bkup -bkup_channels_node=8 -dbname=mini`

The example above will set 8 RMAN channels per node in the cluster up to a maximum of 64 total channels across the cluster.

To change the compression algorithm used by the backup API, set the option bkup_rman_compression. This option accepts the following values: LOW (default), MEDIUM, and High.

Example: `# /var/opt/oracle/ocde/assistants/bkup/bkup -bkup_rman_compression="MEDIUM" -dbname=mini`

The example above will modify the default backup compression algorithm from LOW to MEDIUM.

A sample backup configuration file is shown below:

```
bkup_disk=no
bkup_oss=yes
bkup_oss_url=https://myobjectstorage.uk-london-1.oraclecloud.com/v1/mytenancy/bkup_bucket
bkup_oss_user=ociuser1
bkup_oss_passwd=********
bkup_oss_recovery_window=90
bkup_daily_time=01:00
bkup_cron_entry=yes
bkup_channels_node=8
bkup_rman_compression=MEDIUM
```

Uncompressed backups must be performed outside the backup API. The following steps are required prior to starting the backups.

As user root, comment the backup entries in /etc/crontab referencing the database being backed up.

```
#53 */1 * * * root /var/opt/oracle/bkup_api/bkup_api bkup_archlogs --dbname=mini
#01 01 * * * root /var/opt/oracle/bkup_api/bkup_api bkup_start --dbname=mini
```

As Oracle user, make the modification below in the RMAN configuration.

```
RMAN> CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 8 BACKUP TYPE TO BACKUPSET;
RMAN> CONFIGURE COMPRESSION ALGORITHM CLEAR;
```

The backup, restore and restore validate scripts shown in Appendix A and B can be used to back up the database to the object storage.

After performing the steps above, weekly full backup, daily incremental backups and periodic archive log backups must now be manually created for the database.

ORACLE

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/oracle
facebook.com/oracle
twitter.com/oracle
oracle.com

Integrated Cloud Applications & Platform Services

Oracle Cloud Infrastructure Exadata Backup & Restore Best Practices using Cloud Object Storage
May 2019
Author: Jony Safi
Contributing Authors: Lawrence To; Kelly Smith; Markus Michalewicz