Oracle Maximum
Availability Architecture

Oracle Database Backup Cloud Service

Best Practices for On-Premise Database Backup & Recovery

ORACLE WHITE PAPER | JUNE 2019

Contents

Introduction

Oracle Database Backup Cloud Service is an easy to deploy, secure and scalable subscription service for backing up Oracle on-premise or Cloud databases to the public Cloud. The service complements existing RMAN disk backup strategies by providing a secondary, off-site storage location in the Cloud with unlimited capacity. The service also ensures that backups are encrypted and available when needed.

This service, and not Database Administrators, handle storage management and data transfer complexities. Database Administrators only use the familiar Recovery Manager (RMAN) interface to perform backup and restore operations; no new tools or commands are needed. If the Database Administrator knows how to run RMAN backups to tape or disk, they know how to back up to the Oracle Cloud.

This paper describes the Backup Cloud Service and how it works, along with key best practices for configuration and operational usage for on-premise databases. For best practices related to Cloud databases, refer to MAA Best Practices for Oracle Cloud Backups

Why Backup to the Cloud?

Storing database backups off-site is critical for business continuity in the event of major disasters or outages. Those backups must be accessible 24 x 7 to reduce application downtime.

Off-site backup is traditionally accomplished by sending backups to tape and shipping them to a secure location. This is a complex task that requires hardware, personnel, and procedures to ensure off-site backups are current, validated, and available at a moment's notice.

Oracle Cloud Infrastructure Object Storage provides a great alternative to writing, shipping, and storing tapes at an off-site location which increases performance, redundancy, and security.

Traditional Database Backup Best Practices

The following table summarizes traditional Database Backup Best Practices.

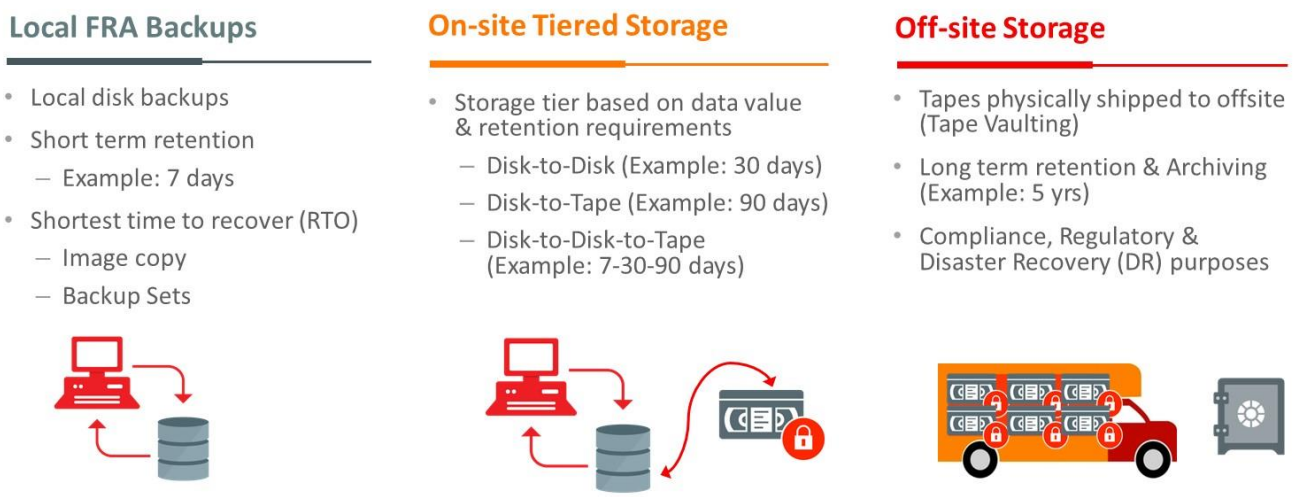

Challenges with Traditional Backup Infrastructure

- **On-Demand Capacity Growth:** With explosion of data growth, storage capacity planning needs to be more agile, especially for long-term retention backup which may be kept for years.
- **Access Delays:** With tape vaulting, offsite data needing to be restored must first be recalled and shipped back to original location - thus, the data is not immediately accessible, increasing overall RTO.
- **High Cost:** Infrastructure and operational expenditures to procure and manage onsite and offsite tape infrastructure continue to rise, as economics of disk becomes more attractive.

Oracle Database Backup Cloud Service: Overview

Disaster can strike without warning. With Oracle Database Backup Cloud Service, your backups are easy to access, secure over the Internet, and are immediately available for recovery when needed.

Oracle Database Backup Cloud Service is simple to deploy and easy to use. Subscribe to the service, install the cloud backup module, configure a few settings, and take your first backup to the cloud using familiar commands and tools. It's that simple.

Oracle Database Backup Cloud Service protects data by providing end-to end security. Data is encrypted at the source, securely transmitted to the cloud, and securely stored in encrypted format. The keys are stored on-site, not in the cloud.

All the backup data stored in the Oracle Database Backup Cloud Service is automatically and transparently replicated across multiple storage nodes in the same geographic region, which provides instant availability.

Setting up Oracle Database Backup Cloud Service in 4 simple steps.

1. Subscribe to the Oracle Database Backup Cloud Service at cloud.oracle.com (or) work with an Oracle representative. Alternatively, click on the "Try for free" button. For more information, refer to cloud.oracle.com/database_backup.

2. Download and install the Oracle Database Cloud Backup Module from Oracle Technology Network (OTN). The module makes it possible to perform secure cloud backups and restores. Install this module on the systems where the Oracle database is running. Multiple database and operating systems versions are supported.

3. Configure RMAN to use the installed module.

4. Start performing backup & recovery operations to the cloud using familiar RMAN commands.

The online dashboard is used to monitor the service and see how much storage capacity is being used for backups.

ORACLE DATABASE BACKUP CLOUD SERVICE – SUPPORTABILITY MATRIX

Database / Features	Supported Versions / Options
Oracle Database – Enterprise Edition*	11.2.0.4, 12c, 18c,19c (64 bits)
Oracle Database – SE/SE1/SE2*	11.2.0.4 and above
Platforms (**64 bit**)	Linux, Solaris, SPARC, Windows, HP-UX, AIX, zLinux
RMAN Compression (Included)	HIGH, MEDIUM, BASIC, LOW**
RMAN Encryption (Included)	Password, TDE, Dual-mode

* Older Database versions no longer supported by Oracle are in deprecated mode

** SE Database supports BASIC only

Oracle Database Backup Cloud Service supports the following RMAN operations:

Database	Backups From Fast Recovery Area	Restore from Cloud	Maintenance
Backup Sets	Image Copies	Full Database	Retention Period
Full Database	Backup Sets	Tablespace	Crosscheck
Selected Tablespace(s)	Archived Logs	Data File	Delete Obsolete
Selected Data Files	Compressed Encrypted	Table Recovery (DB 12c and above)	Delete Backups
Incremental – Differential		Block Recovery	
Incremental – Cumulative			
Compressed Encrypted			

Oracle Database Backup Cloud Service includes RMAN Compression and Encryption:

RMAN Compression (Optional)	RMAN Encryption (Mandatory)

HIGH, BASIC, MEDIUM, LOW
MEDIUM recommended
No ACO licensing required
`CONFIGURE COMPRESSION ALGORITHM 'MEDIUM';`
`BACKUP AS COMPRESSED BACKUPSET DATABASE PLUS ARCHIVELOG;`
	Password, Transparent Data Encryption (TDE), Dual-Mode
	No ASO licensing required
	Keys are stored locally (not in Cloud Storage)
	If TDE is used (preferred), then simply use `SET ENCRYPTION ON` before backups and restores
	For password encryption:
	`SET ENCRYPTION ON IDENTIFIED BY '<password>' ONLY;`
	Before doing restore,
	`SET DECRYPTION IDENTIFIED BY '<password>';`

About the Database Backup Cloud Module Install

The Database Backup Cloud Service supports both the legacy Oracle Cloud Infrastructure Classic Object Storage (OCI-C) and new Oracle Cloud Infrastructure Object Storage (OCI).

The Database Cloud Backup Module comes with two different Java installer modules:

- `opc_install.jar` is the installer module to set up backups to OCI-C
- `oci_install.jar` is the installer module to set up backups to OCI.

If `opc_installer` was previously used to set up backups to OCI-C or OCI (via Swift compatible endpoints), it is highly encouraged to switch to the new module by running the `oci_installer`. See Migrating from OCI-C containers to OCI buckets below for more details.

INSTALLATION OPTIONS - SUMMARY

Target Cloud Infrastructure	Installer Module	Endpoint format	Notes
Oracle Cloud Infrastructure Classic	`opc_install,jar`	`https://<domain>.oraclecloud.com/v1/Storage-<domain>`	This is the only option when backing up to the OCI Classic Object Storage.
Oracle Cloud Infrastructure Native APIs	`oci_install.jar`	`https://objectstorage.<region>.oraclecloud.com`	This is the recommended approach for OCI. It uses a stronger key based authentication scheme and will support future OCI features.
Oracle Cloud Infrastructure via Swift APIs	`opc_install.jar`	`https://swiftobjectstorage.<region>.oraclecloud.com/v1/<tenant>`	No longer recommended - This method installs the Swift API based OCI Classic module to backup to OCI via the Swift compatible end-point. The OCI native installer above should be used with Oracle Cloud Infrastructure.

Note: This White Paper refers only to the new OCI native module to backup to OCI Object Storage buckets. Please refer to the DB Backup Cloud Service documentationD for OCI Classic information.

HIGH LEVEL ARCHITECTURE OF ORACLE CLOUD BACKUP

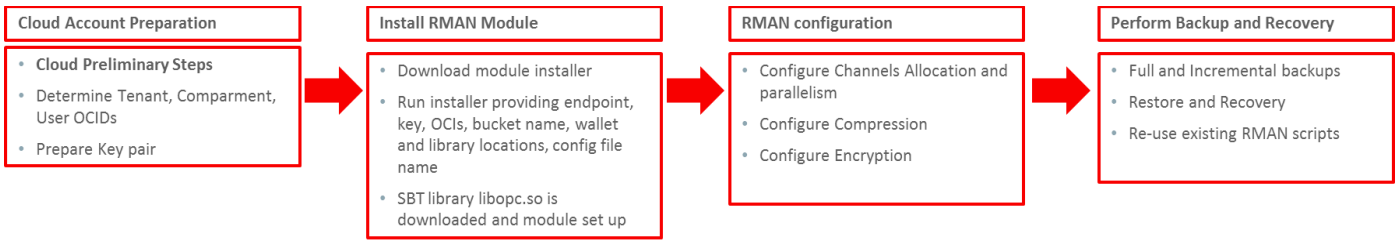

Database Cloud Backup Module

The Database Cloud Backup Module is a system backup to tape (SBT) interface that is tightly integrated with Recovery Manager (RMAN), which means you don't need to learn new tools or commands.

You can continue to use standard RMAN commands for all backup, restore, recovery, and maintenance operations.

Download the backup module from Oracle Technology Network (OTN) and install it on your database server. Multiple database versions and operating systems are supported. For more information about the module, see Installing the Oracle Database Cloud Backup Module.

Workflow:

1. RMAN reads backup data from the database and sends it to the Oracle Cloud Backup module.
2. The Cloud module breaks backup pieces into 100MB chunks (Default) and sends them to the Cloud.
    a. Any failed transmissions are retried automatically.
    b. Multiple buffers (RMAN Channels) can be used for parallelism and to increase backup throughput if there is sufficient internet network bandwidth.
3. Each chunk is stored as an object inside the Oracle Cloud bucket. The buckets can either be user pre-created (or) automatically created by the Cloud Backup module. The default bucket created by the RMAN cloud module will be named: "oracle-data-[first 8 chars of service & tenant]"
4. REST API calls – `PUT,GET, POST, HEAD, & DELETE` are used over HTTPS.
5. Typical URL formation for every object.
    a. `https://objectstorage.<region>.oraclecloud.com/n/tenant/<bucket>/<piece name>/<unique ID>/0000001, 0000002 ..`
6. An XML manifest file is created and maintains metadata for the chunk files in the Cloud, which is used by the Cloud Backup module.

Database Cloud Backup Module Files

File name	Location / Creation	Purpose
`libopc.so(or) oraopc.dll`	User specified library location. Downloaded by the installer.	SBT library which enables backup to Oracle Cloud
`opc<SID>.ora`	Configured by the installer under `$ORACLE_HOME/dbs` or user specified location	Contains bucket name URL location for the user and also the credential wallet location
`cwallet.sso`	User specified wallet location created during the RMAN module installation.	Oracle wallet which securely stores backup service credentials. This is used during RMAN backups and restore operations .
Wallet for encryption (optional –only needed for TDE)	Either `$ORACLE_BASE /admin/$ORCLE_SID /wallet` (or) defined in `sqlnet.ora` / Existing wallet	Used for backup encryption. Existing Oracle wallet can be used (or) new Oracle wallet can be created.

Preparing to Run the Database Cloud Backup Module installer

Before running the installer, you need to gather some information from your cloud account. Follow these steps:

1. Identify your tenant's OCID. You can find it by clicking on the Profile icon on the top right corner of the cloud console and selecting Tenancy: <your tenancy name> from the drop down menu. You will be taken to your tenancy detail information where you can find its OCID. Copy the OCID string and save it in a temporary .location for later use. The OCID string is similar to: `ocid1.tenancy.oc1..aaaaaaaaj62uff362gve2deswibx3tgsgv2ng7nny7fwhz6ecnjdcupor3yq`

2. Identify the compartment where your backup buckets will be placed. An existing compartment can be used or create a new one from the Compartments page. Compartments page can be reached by selecting Identity --> Compartments on the left side of the console menu. Copy the OCID of the existing compartment or create a new one first and save the compartment OCID in a temporary .location for later use. If you don't want to use compartments and prefer to use the root compartment, this step is not necessary.

3. Identify the cloud user that will be responsible for managing the cloud bucket, you can create a new user or use an existing one. The user must have permission to manage buckets and objects in the compartment previously identified. On the user management page (Identity --> Users), copy the OCID for the specific user and save it in a temporary location for later use.

4. Prepare your key pair for API signing in pem format as described here https://docs.cloud.oracle.com/iaas/Content/API/Concepts/apisigningkey.htm. Do not use passphrase protection on the private key. Save your private key file and copy your public key to the user management console page as shown on the documentation. Note itss fingerprint and copy it in a temporary location for later use.

5. Identify the Object Storage API endpoint for the region your backups will be sent to. The endpoint format will be https://objectstorage.<region>.oraclecloud.com. Examples of currently available endpoints are listed below.

    - https://objectstorage.ca-toronto-1.oraclecloud.com
    - https://objectstorage.eu-frankfurt-1.oraclecloud.com
    - https://objectstorage.uk-london-1.oraclecloud.com
    - https://objectstorage.us-ashburn-1.oraclecloud.com
    - https://objectstorage.us-phoenix-1.oraclecloud.com

Running the DB Cloud Backup Module Installer

Run the installer using the parameters that you prepared in advance from above:

`java -jar oci_install.jar -host` <endpoint from #5>

`-pvtKeyFile` <local location of the file containing the private key in pem format from step #4>

`-pubFingerPrint` <public key fingerprint from step #4>

`-tOCID` <tenancy OCID from step #1>

`-cOCID` <compartment OCID from step #2>

`-uOCID` <user OCID from step #3>

`-walletDir` <directory where the installation will store the credential wallet>

`-libDir` <directory where the installation will store the SBT library>

`-configfile` <configuration file name created during installation>

`-bucket` <bucket name from step #2>

Install example:

```
java -jar oci_install.jar \
-host https://objectstorage.us-ashburn-1.oraclecloud.com \
-pvtKeyFile ~/oci_api_key.pem \
-pubFingerPrint 21:b1:ab:a0:b0:f0:50:30:ee:d6:a7:18:b3:50:a8:36 \
-tOCID ocid1.tenancy.oc1..aaaaaaaaj4ccqe763dizkrcdbssx7ufvlmokd24mb6utvkymyo2xwxyv3gfa \
-cOCID ocid1.compartment.oc1..aaaaaaaaxslr7vtt5cj4ksb3lvwu6agbvo5gh7t5iljd4ydfolgfy4wdpnrq \
-uOCID ocid1.user.oc1..aaaaaaaaid4hi2kzgbbyzjtietoaxxh2gzk4r2bqqqxwag7cqli5cpw6ls4a \
-walletDir ~/ociwallet -libDir ~/ocilib -configfile ~/ociconfig/opcORCL.ora \
-bucket OCIbucket
```

For more details, refer to the DB Backup Service documentation:
https://docs.oracle.com/en/cloud/paas/db-backup-cloud/csdbb/installing-oracle-database-cloud-backup-module-oci.html

CONTENT OF THE OPC CONFIGURATION FILE - RUN TIME CONFIGURATIONS*

Parameter Name	Description
OPC_HOST	REST destination URL Ex: https://objectstorage.<region>.oraclecloud.com/n/<tenant>
OPC_WALLET	OPC credential wallet location Ex: 'LOCATION=file:/home/oracle/OPC/wallet CREDENTIAL_ALIAS=odbs_opc'
OPC_CONTAINER	User specified bucket name Ex: PAYROLL_DB (can be created using the Object Storage console)
OPC_CHUNK_SIZE	Not recommended to change. Size of stored backup checks. specified in bytes. By default, 100MB.
OPC_COMPARMENT_ID	Target compartment OCID
_OPC_TRACE_LEVEL	For debug purposes only. Set this parameter to – say 100 which generates more trace information in sbtio.log.

* Default location: $ORACLE_HOME/dbs/opc<sid>.ora

ARCHITECTURE OF ORACLE CLOUD BACKUP WHEN USED WITH ON-PREMISE DATABASES

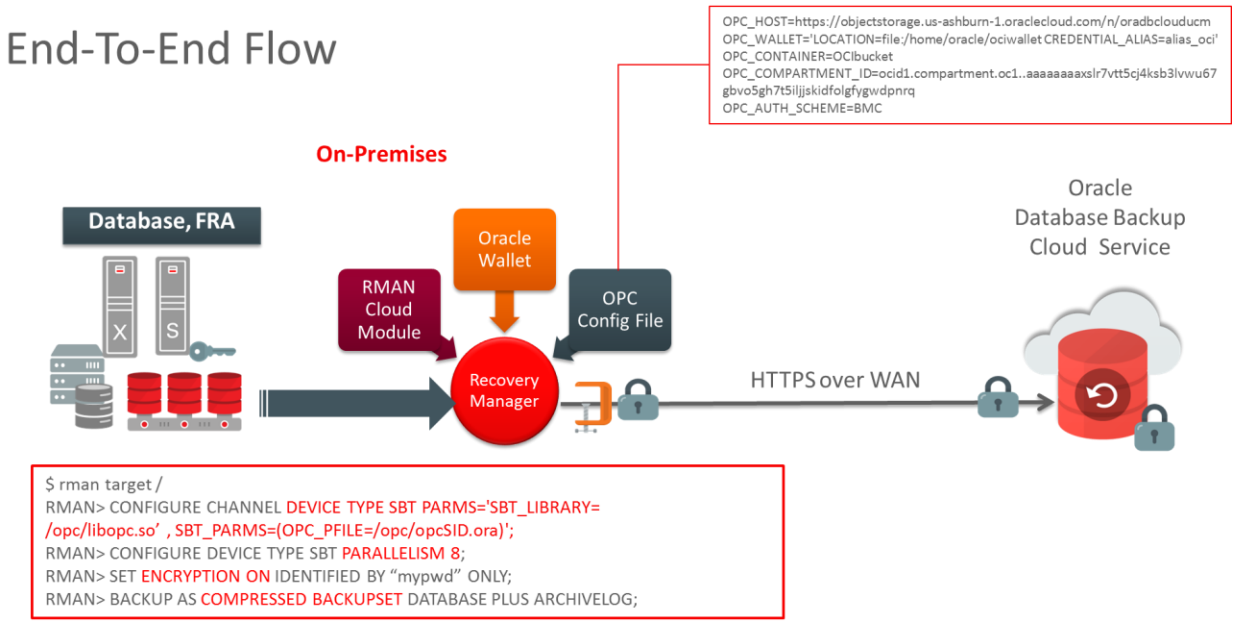

Oracle Database Backup Cloud Service: RMAN Best Practices

This section discusses the best practices for backing-up or recovering from the Oracle Cloud Backup Service. These are based on native RMAN commands.

Before starting, ensure the Oracle Cloud Backup module from OTN has been installed and the RMAN environment has been configured properly.

```
RMAN>CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS
'SBT_LIBRARY=/home/oracle/OPC/lib/libopc.so,
ENV=(OPC_PFILE=/u01/products/db/12.1/dbs/opcodbs.ora)';
```

Backup Best Practices

- Use RMAN encryption for backups. This is enforced and mandatory when backing up On-Premise Databases. The RMAN set encryption clause in your RMAN run block will ensure that encryption is enabled.

```
RMAN> SET ENCRYPTION ON IDENTIFIED BY 'abc123' ONLY;
```

  Keys are managed by the customer (password, TDE, dual-mode) and data will be securely transmitted to the cloud over HTTPS

- Use compression and parallelism to optimize data transfer when network bandwidth is limited and CPU resources are available.
  - RMAN compression (`HIGH, MEDIUM, LOW, BASIC`)

```
RMAN> CONFIGURE COMPRESSION ALGORITHM 'MEDIUM';
```

```
RMAN>BACKUP DEVICE TYPE SBT AS COMPRESSED BACKUPSET DATABASE PLUS
ARCHIVELOG FORMAT '%d_%U';
```

  - Parallelism can be increased until acceptable network throughput or the maximum internet throughput is reached.

```
RMAN> CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 4 BACKUP TYPE TO
BACKUPSET;
```

To determine network throughput for a specific time period, use RMAN network analyzer, see MOS note **2022086.1**

To diagnose Oracle Cloud Backup Performance, see MOS note **2078576.1**.

- Use `MULTISECTION` backups

  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 file in separate sections can improve the performance of backups of large data files. For example, suppose that the "users" tablespace contains a single datafile of 800 MB and assume that four SBT channels are configured, with the parallelism setting for the SBT device set to 4. The example shown below breaks up the datafile in the "users" tablespace into 4 sections, which are backed up in parallel across the 4 channels.

```
RMAN> BACKUP SECTION SIZE 200M TABLESPACE USERS;
```

- Use "weekly full and daily incremental" strategy

  The goal of an incremental backup is to back up only those data blocks that have changed since a previous full or incremental backup.

  The advantages of this strategy are:

  - Reduces the amount of time needed for daily backups, as only changed blocks are backed up. Incrementals may be taken more frequently (e.g. twice a day) to further reduce RPO.
  - Reduces network usage and network bandwidth requirements when backing up over an internet network.
  - Reduces backup overhead and read I/O, with RMAN block change tracking feature.

  The tradeoff with incrementals is that the recovery time can take longer as incremental backups must be restored and applied, after data files are restored.

Example of a Weekly full/daily incremental strategy:

Sunday

An incremental level 0 (Full) backup saves all blocks in the database.

```
RMAN> BACKUP INCREMENTAL LEVEL 0 DATABASE PLUS ARCHIVELOG NOT
BACKED UP DELETE INPUT;
```

Monday - Saturday

On each day from Monday through Saturday, a differential incremental level 1 (Incremental) backup saves all blocks that have changed since the most recent backup at level 1 or 0. So, the Monday backup saves blocks changed since Sunday level 0 backup, the Tuesday backup saves blocks changed since the Monday level 1 backup, and so forth.

```
RMAN> BACKUP INCREMENTAL LEVEL 1 DATABASE PLUS ARCHIVELOG NOT
BACKED UP DELETE INPUT;
```

The RMAN block change tracking feature for incremental backups improves incremental backup performance by recording changed blocks in each data file in a change tracking file. If change tracking is enabled, RMAN uses the change tracking file to identify changed blocks for incremental backup, thus avoiding the need to scan every block in the data file at backup time.

To enable or disable block change tracking refer to the example below. Additional information can also be found in the RMAN Incremental Backup documentation.

```
SQL>ALTER DATABASE ENABLE BLOCK CHANGE TRACKING;
SQL>ALTER DATABASE DISABLE BLOCK CHANGE TRACKING;
```

In summary, the RMAN configuration should contain similar settings to those shown below:

```
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS
'SBT_LIBRARY=/home/oracle/OPC/lib/libopc.so,
ENV=(OPC_PFILE=/u01/products/db/12.1/dbs/opcodbs.ora)'
CONFIGURE COMPRESSION ALGORITHM 'MEDIUM’
CONFIGURE CONTROLFILE AUTOBACKUP ON
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 4 BACKUP TYPE TO
BACKUPSET
CONFIGURE BACKUP OPTIMIZATION ON
```

After the backup is complete, the backups can be displayed by using the RMAN list command. Note: the Media attribute name refers to the location in the Oracle Cloud Service.

```
RMAN> LIST BACKUP;


BS Key  Type LV Size       Device Type Elapsed Time Completion Time
------- ---- -- ---------- ----------- ------------ ---------------
714     Full    Unknown    SBT_TAPE    00:00:10     29-MAR-19
```

```
          BP Key: 787   Status: AVAILABLE  Compressed: YES  Tag:
TAG20190329T224129
          Handle: ORCL_1527520098_mbttme1q_1_1_20190329_1004222522
Media: objectstorage.us-ashburn-1..ecloud.com/n/oratenant/OCIbucket
  List of Datafiles in backup set 714
  File LV Type Ckp SCN    Ckp Time  Abs Fuz SCN Sparse Name
  ---- -- ---- ---------- --------- ----------- ------ ----
  5       Full 1230808    09-FEB-19             NO
+DATA/ORCL_IAD1D2/72C8DB3ED2DD02D9E053060011AC9203/DATAFILE/system.266.
999739411.
.
.
```

Recovery Best Practices

Because accidents can happen and often without warning, you need to ensure that your backups are available when you need them. Oracle Cloud Backup offers you performance, redundancy, and security, which in turn provide peace of mind. Nevertheless, proactively testing restore and recovery procedures is still an important activity and should be conducted regularly.

Recovery is commonly required in the event of:

- Storage Failure
- Block Corruption
- User/Logical Error
- Database Failure
- Site failure or disaster

Consult the following Database MAA best practices to detect, prevent, and repair from data corruptions.

- Preventing, Detecting, and Repairing Block Corruption: Oracle Database 12c
- Preventing, Detecting, and Repairing Block Corruption: Oracle Database 11g

Cross-check Backups Best Practices

Cross-checking backups is important and should be done before a `delete obsolete` command. Cross-checking marks the missing backup set/piece as expired in the RMAN repository (control file and/or RMAN catalog) and does not delete or remove the actual files. Backup set/pieces marked as expired are excluded from subsequent `backup, restore, recover, delete obsolete` commands.

Following a `crosscheck` command, it is recommended to run `report expired` to review and confirm any missing backup files. Then run `delete expired` to remove the entries flagged as expired from the RMAN repository.

Use `crosscheck` to check that files are accessible and ready for a restore operation.

```
RMAN> CROSSCHECK BACKUP;
RMAN> CROSSCHECK BACKUP OF DATABASE;
RMAN> LIST EXPIRED BACKUP OF DATABASE;
```

Validate Backups Best Practices

As storage media can become corrupted for various reasons, RMAN provides mechanisms to check for physical and logical corruption on backups.

RMAN `restore validate` command does a block level check of the backups and verifies all needed database files are available, thus ensuring that an actual restore can be performed. It is recommended to validate backups on a regular basis.

```
RMAN> RESTORE DATABASE VALIDATE CHECK LOGICAL;
```

Note: RMAN `restore validate` reads the backup sets and checks them for corruption. RMAN `restore validate` does consume CPU, memory and network resources to read the backups and analyze them. However, no data is written to storage. The data is streamed from the cloud to your on-premises database for validation purposes and is discarded after the validation. You may incur network traffic charges for data leaving the Oracle Cloud after the 10TB/month free tier.

For large backup sets, `restore validate` command can take longer to complete. For a quick validation to ensure the backup files are available you can leverage `restore validate` **`header`**. This will validate that backups are present but will not perform block-level check.

```
RMAN>RESTORE DATABASE VALIDATE HEADER;
```

Use `backup validate` after a backup completes to validate the database data files. The `validate` clause will check for physical corruption only in used blocks. To extend the check for logical corruptions, use `check logical` in conjunction with the `validate` clause.

```
RMAN>BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG
ALL;
```

In Summary:

- `Crosscheck`: Ensures that the backup pieces are available on the Cloud object store. It compares the backup metadata (either in the controlfile or catalog) against the physical backup pieces to check if it matches.
- `Backup validate`: Checks the database data files for physical corruptions. With the `check logical` option, the command checks for logical corruptions as well.
- `Restore validate`: Checks if the backup is restorable and if it contains any physical corruptions and with the check logical option, the command checks for logical corruptions as well.

Example Plan:

1. Daily RMAN `crosscheck`: To ensure that backup pieces are available for restore.

2. Monthly `restore validate` with `check logical`: To confirm that a restore can be performed in the event of a disaster.

3. Quarterly Full Restore and Recovery: To test DR strategy.

Additional Best Practices:

- Use RMAN `LOW` or `MEDIUM` compression for optimal data transfers
- Increase PARALLELISM (until maximum network throughput is reached)
- Refer to MOS Note 2078576.1 for performance investigation of backups
- If public network throughput is not sufficient, choose Oracle Fast Connect (Standard, Partner Edition, MPLS). Refer to cloud.oracle.com/networking
- Choose cloud standard or archive storage as appropriate based on RTO/RPO
- Perform weekly full and daily incremental backups
- Schedule archived logs backup frequency to reduce RPO
- Run Installer once every two months to update to the latest RMAN SBT module
- Copy `opc<SID>.ora` file to other SIDs if same ORACLE_HOME is used by multiple databases
- Configure `CONTROLFILE AUTOBACKUP ON`. This will enable complete restore of a database into a different host.

Updating from the Swift-based Legacy Module to the OCI Native Module

If OCI Object Storage is being used with the legacy module via the Swift endpoints, all that is needed is to download the new Database Cloud Backup Module. Then run the `oci_install.jar` installer specifying the Object Storage endpoint for the region with the appropriate authentication parameters and the existing bucket name. Existing backups will now be accessed using the OCI native APIs. No other action is required. Although the RMAN catalog will continue to show the Swift endpoint in the "Media:" field for backup pieces created by the legacy module, this is just a label and can be ignored.

Migrating Backups from OCI-C Object Storage Classic to OCI Object Storage

If on-premises databases are being backed up to OCI-C Object Storage Classic and migration to OCI Object Storage is required, decide if existing backups need to be moved. There are two approaches based on the retention requirements of the backups already created.

If the recovery window is short, just install the new OCI native backup module and start backing up to an OCI bucket. Make sure to specify a different location for the credentials wallet and a different opc configuration file. The same location is kept for the `libopc.so` SBT library, as the library itself is not different. Doing this will start fresh backups on a new OCI bucket. If a restore is needed from backups taken with the legacy module, still located in the OCI-C container, just use the previous configuration file in the channel allocation parameter. The backups can be read from their original location. As the recovery window slides forward over the coming days, the old backups will all eventually become obsolete and all recent backups will be in the new OCI bucket.

If the recovery window is long and there are backups that need to be kept for a long time, it would be best to copy them from the OCI-C container to the OCI bucket in order to retain access to them.

1. Prepare the OCI target destination (user, compartment, bucket)

2. Use a tool like `rclone` to copy the whole content of the OCI-C container to the new OCI bucket. The process is described in the White Paper "Transferring Data to Object Storage from Other Cloud Providers or Local File Systems" available here: https://cloud.oracle.com/iaas/whitepapers/transfer-data-to-object-storage.pdf

   Below is an example of `rclone` settings used to migrate backups from OCI-C container OPCbucket to OCI bucket OCIbucket:

   Source:

   OCI-C domain id: domain123
   OCI-C container name: OPCbucket
   OCI-C user: user1@mycompany.com
   OCI-C password: MyPassword

Destination:

OCI region: us-ashburn-1 region
OCI tenancy:mytenancy
OCI user and authentication key specified as S3 ID and Access Key

```
export RCLONE_CONFIG_OCIC_TYPE=swift
export RCLONE_CONFIG_OCIC_USER=Storage-domain123:user1@mycompany.com
export RCLONE_CONFIG_OCIC_KEY=MyPassword
export RCLONE_CONFIG_OCIC_AUTH=https://Storage-domain123.storage.oraclecloud.com/auth/v1.0
export RCLONE_CONFIG_OCIC_AUTH=https://uscom-east-1.storage.oraclecloud.com/auth/v1.0
export SOURCE=ocic:OPCbucket
export RCLONE_CONFIG_OCI_TYPE=s3
export RCLONE_CONFIG_OCI_ACCESS_KEY_ID=dcc9f5358c1479081442e7cdbf6ca72836fe9
export RCLONE_CONFIG_OCI_SECRET_ACCESS_KEY=pcBXigCzxzfeDeoFC8EVrLBJd0B/g+v4m3co
export RCLONE_CONFIG_OCI_REGION=us-ashburn-1
export RCLONE_CONFIG_OCI_ENDPOINT=https://mytenancy.compat.objectstorage.us-ashburn-
1.oraclecloud.com
```

Once these variables are set the following command will copy all the content from OCI-C OPCbucket to OCI OCIbucket

```
rclone --verbose --cache-workers 64 --transfers 64 --retries 32 copy $SOURCE oci:OCIbucket
```

3. Download the new Database Backup Cloud Module and run the `oci_install.jar` installer pointing to the destination bucket.

4. Perform a `restore validate` to verify your backups are still accessible.

Conclusion

Businesses are increasingly evaluating and moving their on-premise environments to the Cloud for lower cost, easier management, and unlimited scale. Backups are commonly viewed as initial candidates for moving to Cloud, due to the cost of managing traditional tape backup and vaulting infrastructure, including cost of maintaining an offsite backup location.

With Oracle Database Backup Cloud Service, customers now have an effective and low cost solution to protect their Oracle databases along with storing backups in an offsite, secure, and anytime-anywhere accessible location. Configuration and operational best practices detailed in this paper will further ensure that backups to and recovery from the Cloud are best optimized for your on-premise and Cloud database environments.

References

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

DB Backup Cloud Service on cloud.oracle.com

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

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US

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

Integrated Cloud Applications & Platform Services

Oracle Database Backup Cloud Service
Backing up Oracle Databases to Oracle Cloud Infrastructure

June 2019