



Oracle Database Backup Cloud Service



Best Practices for On-Premises Database Backup & Recovery.

April, 2021 | Version 05.40

Copyright © 2021, Oracle and/or its affiliates

Confidential - Public/Oracle Internal/Oracle Restricted/Oracle Highly Restricted

PURPOSE STATEMENT

This document provides guidance and best practices for Backup and Restore of on-premises Oracle Databases to and from the Oracle Database Backup Cloud Service.

Oracle Database Backup Cloud Service is an easy to deploy, secure and scalable service for backing up Oracle on-premises or Cloud databases to the public Cloud. The service can be the primary backup destination or complement existing RMAN disk backup strategies by providing a virtually unlimited, off-site storage location in the Cloud. The service also ensures that backups are encrypted and available when needed.

INTENDED AUDIENCE

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

DISCLAIMER

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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 in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

TABLE OF CONTENTS

PURPOSE STATEMENT	1
INTENDED AUDIENCE	1
DISCLAIMER	1
INTRODUCTION	3
EXECUTIVE SUMMARY	3
TRADITIONAL DATABASE BACKUP BEST PRACTICES	3
CHALLENGES WITH TRADITIONAL BACKUP INFRASTRUCTURE	4
ORACLE DATABASE BACKUP CLOUD SERVICE: OVERVIEW	4
SETTING UP ORACLE DATABASE BACKUP CLOUD SERVICE	5
ABOUT THE DATABASE BACKUP CLOUD MODULE INSTALLER	7
4 STEPS TO CLOUD BACKUPS	8
DATABASE CLOUD BACKUP MODULE WORKFLOW AND COMPONENTS	8
OBJECT LIFECYCLE POLICIES AND STORAGE TIERS	9
PREPARING TO RUN THE DATABASE CLOUD BACKUP MODULE INSTALLER	10
RUNNING THE DB CLOUD BACKUP MODULE INSTALLER	11
ORACLE DATABASE BACKUP CLOUD SERVICE: RMAN BEST PRACTICES	14
BACKUP BEST PRACTICES	14
RECOVERY BEST PRACTICES	17
CROSS-CHECK BACKUPS BEST PRACTICES	17
VALIDATE BACKUPS BEST PRACTICES	18
VALIDATION BEST PRACTICES SUMMARY	18
ADDITIONAL BEST PRACTICES	19
CONCLUSION	19
APPENDIXES	20
USING OBJECT STORAGE REPLICATION	20
UPDATING FROM THE SWIFT-BASED LEGACY MODULE TO THE OCI NATIVE MODULE	22
MIGRATING BACKUPS FROM OCI-C OBJECT STORAGE CLASSIC TO OCI OBJECT STORAGE	22
REFERENCES	23

INTRODUCTION

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.

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 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.


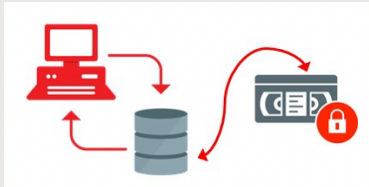

EXECUTIVE SUMMARY

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.

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.

TRADITIONAL DATABASE BACKUP BEST PRACTICES

The following table summarizes traditional Database Backup Best Practices.

LOCAL FRA BACKUPS	ON-SITE TIERED STORAGE	OFF-SITE STORAGE
<ul style="list-style-type: none"> Local Disk Backups Short Term Retention (Example: 7 Days) Short Time to Recover (RTO) Image Copy Backup Sets Limited HW Failure Protection 	<ul style="list-style-type: none"> Storage Tier based on Business Value and Retention Requirements Disk-To-Disk (Example: 30 Days) Disk-To-Tape (Example: 90 Days) Disk-To-Disk-To-Tape (Example: 7-30-90 Days) 	<ul style="list-style-type: none"> Tapes physically shipped Off-Site Long Term Retention and Disaster Recovery Compliance Archiving (Example: 5-10 Years) 

CHALLENGES WITH TRADITIONAL BACKUP INFRASTRUCTURE

Traditional backup solutions, like those described in the table above have some inherent challenges:

- **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

Oracle Database Backup Cloud Service leverage the Oracle Cloud Infrastructure Object Storage service to provide, virtually unlimited, offsite backup storage. There are no more scalability issues and backups are stored offsite in a durable, secure location.

No need for complicated capacity planning exercises for hardware purchases. The Object Storage service is always available, and you pay for what you use.

Oracle Database Backup Cloud Service protects data by providing end-to end security. Data are 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. Optional cross-region replication can be set up at the bucket level using OCI Object Storage Replication.

Oracle Database Backup Cloud Service – Supportability Matrix

DATABASE/FEATURES	SUPPORTED VERSIONS
Oracle Database – Enterprise Edition*	Any Supported Database Version (64 bits)
Oracle Database – SE/SE1/SE2*	Any supported Database Version (64 bits)
Platforms (64 bits)	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

The Oracle Database Backup Cloud Service includes the ability to use the RMAN Encryption (mandatory) and Advanced Compression that would otherwise require the licensing of the Advanced Security and Advanced Compression Database Options.

RMAN ENCRYPTION (MANDATORY)	RMAN COMPRESSION (OPTIONAL)
Password, Transparent Encryption, Dual-Mode <u>No ASO licensing required</u> Keys are stored locally in the DB encryption wallet (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>' ;	HIGH, BASIC, MEDIUM, LOW MEDIUM or LOW recommended <u>No ACO licensing required</u> CONFIGURE COMPRESSION ALGORITHM 'MEDIUM' ; BACKUP AS COMPRESSED BACKUPSET DATABASE PLUS ARCHIVELOG ;

SETTING UP ORACLE DATABASE BACKUP CLOUD SERVICE

Starting backing up to the cloud is very easy in just 4 steps.

1. Subscribe to the Oracle Cloud Infrastructure at cloud.oracle.com (or) work with an Oracle representative. Alternatively, you can try the service for free. For more information, refer to <https://www.oracle.com/cloud/free/>
2. Download and install the Oracle Database Cloud Backup Module from [Oracle.com](https://www.oracle.com). 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.

There is no learning curve in order to start using the DB Backup Cloud Module. All the backup and restore operations are performed using the familiar interface of Recovery Manager.

Oracle Database Backup Cloud Service supports the following RMAN operations:

DATABASE	BACKUP RECOVERY AREA, BACKUP BACKUPSET	RESTORE FROM CLOUD	MAINTENANCE
Backup Sets	Image Copies	Full Database	Retention Period
Full Database	Backup Sets	Tablespace	Crosscheck
Selected Tablespace(s)	Archived Logs	Datafile	Delete Obsolete
Selected Datafile(s)	Compressed Backups Encrypted Backups	Table Recovery (12c and above)	Delete Backups
Incremental Differential		Block Recovery	Encrypted Backups (Password, Transparent, and Dual Mode)
Incremental Cumulative			
Compressed Backups Encrypted Backups			



ABOUT THE DATABASE BACKUP CLOUD MODULE INSTALLER

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

- `oci_installer.jar` is the installer module to set up backups to OCI using the OCI native APIs.
- `opc_installer.jar` is the installer module to set up backups to OCI using the Swift-compatible APIs.


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 OCI native APIs by running the `oci_installer`. See [Migrating from OCI-C containers to OCI buckets](#) below for more details.

NOTE: This paper only refers to the OCI native module to backup to OCI Object Storage buckets.

4 STEPS TO CLOUD BACKUPS

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.com](https://www.oracle.com) 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](#).



DATABASE CLOUD BACKUP MODULE WORKFLOW AND COMPONENTS

This is the workflow that happens when the database sends backup to the cloud using the Cloud Backup module:

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 in a multi-part upload operation. Chunks are re-combined into a single object by the Object Storage service when the multi-upload operation is committed. As up to 10,000 chunks are used per multi-part operation, any backup piece with a size of 1TB (100MB X 10,000) or less will be represented by a single data object in the bucket.
3. Any failed chunk transmission is retried automatically.
4. Multiple buffers (RMAN Channels) can be used for parallelism and to increase backup throughput if there is sufficient network bandwidth.
5. Each backup piece up to 1TB in size is stored as an object inside the Oracle Cloud bucket, larger backup pieces will create multiple objects. The buckets can either be user pre-created (or) automatically created by the Cloud Backup module installer.
6. REST API calls – PUT, GET, POST, HEAD, & DELETE are used over HTTPS. So the only IP Port that needs to be opened between the DB server and the OCI endpoint is https (443 TCP port). The DB Cloud Backup Module supports the use of a proxy server.
7. Typical OCI URL format for every object.

```
https://objectstorage.<region>.oraclecloud.com/n/<storage_namespace>/b/<bucket>/o/<piece name>/<unique ID>/0000001, 0000002 ..
```

8. An XML manifest file is created and maintains metadata for the backup pieces in the Cloud, which is used by the Cloud Backup module.

Oracle Database Cloud Backup Module files:

FILENAME	LOCATION	PURPOSE
libopc.so (or) oraopc.dll	User specified library location. Downloaded by the installer from OCI.	SBT library which enables backup to Oracle Cloud
opc<SID>.ora	Configured by the installer under \$ORACLE_HOME/dbs or user specified location. The name can also be specified when running the installer.	Contains: <ul style="list-style-type: none"> • URL for Object Storage Service endpoint • Bucket Name • 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 keys (only needed for Transparent Mode)	Either \$ORACLE_BASE /admin/\$ORACLE_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.

OBJECT LIFECYCLE POLICIES AND STORAGE TIERS

Once data is stored in the cloud, OCI Object Storage offers three different storage tiers, Standard Storage, Archive Storage and Infrequent Access. Let's examine each one in detail and see how they relate to the DB Backup Cloud Service

- **Standard Object Storage:** this is the normal Object Storage service, which should be used for “hot” data that must be quickly available for retrieval at any time. A backup stored in Standard Object Storage can be immediately restored by RMAN without any delay
- **Archive Storage:** Data in the Standard Object Storage tier can be moved to the Archive Storage tier after a defined period (from 0 days to years) using Object Lifecycle Policies (OLP). The DB Cloud Backup Module installer provides the ability to add an OLP to the bucket where backups will be stored. Data in the Archive tier must be “restored” to the Standard tier before it can actually be retrieved from the bucket. This restore operation requires at least 1 hour and depends on the size of the objects to be restored. RMAN can control this operation, but the delay must be taken into account when considering the Recovery Time Objective. Data stored in the Archive Storage tier are charged at a lower rate, so this tier is particularly effective for compliance archiving backups that must be retained for very long time and have a low probability of being restored. There is a 90-day minimum retention time in Archive Storage. Data stored in Archive tier will be charged for 90 days even if they are deleted before that time. The DB Cloud Backup Module installer can set up the appropriate Lifecycle Policy to enable tiering to Archive Storage.
- **Infrequent Access:** This OCI tier is similar to the Standard Object Storage tier, so data are always immediately available although there is 31 days minimum retention time. It is charged at a lower rate than the Standard tier (but higher than the Archive tier), but there is additional cost for downloaded gigabyte. The Infrequent Access tier is enabled in the OCI Console on a per-bucket basis using Object Lifecycle Policies.

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 new one created from the Compartments page. The 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 a 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 its fingerprint and copy it to a temporary location for later use.
5. If you plan on using Tiering to Archive Storage, determine for how long you want your data to remain in Standard Object Storage before being moved to Archive Storage and also for how long you want your data to remain in Standard Object Storage when a restore from Archive operation is performed before being sent to the archive tier again. The default values are: archive after 0 days (i.e., less than 24 hours) and retain restored objects in Standard Storage for 24 hours before re-archiving them.
6. Identify the Object Storage API endpoint for the region your backups will be sent to. The endpoint format will be in the form of <https://objectstorage.<region>.oraclecloud.com>. Some 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, it must exist>
-libDir <directory where the installation will store the SBT library>
-configfile <configuration file name created during installation>
-bucket <bucket name, if it does not exist it will be created>
[ -enableArchiving <enable tiering to archive storage>
-archiveAfterBackup <time the backups are retained in standard storage the default is : "0 DAYS">
-retainAfterRestore <time backups will remain in standard storage after a restore from archive, the default
is: "24 HOURS"> ]
```

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..aaaaaaaaj4ccqe763dizkrdbssx7ufvokd24mb6utvkymyo2xwxyv3gfa \
-cOCID ocid1.compartment.oc1..aaaaaaaaxslr7vtt5cj4ksb3lvwu6avo5gh7t5iljd4ydfolgyfy4wdpnrq \
-uOCID ocid1.user.oc1..aaaaaaaaid4hi2kzgbbyzjtietoaxxh2gzk4r2bqqqxwag7cqli5cpw6ls4a \
-walletDir ~/ociwallet -libDir ~/ocilib -configfile ~/ociconfig/opcORCL.ora \
```



```
-bucket OCIBucket  
-enableArchive TRUE  
-archiveAfterBackup "2 WEEKS"  
-retainAfterRestore "48 HOURS"
```

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>

NOTE: Using tiering to Archive Storage is achieved through the use of Object Lifecycle Policies and requires that the Object Storage Service in the region is authorized to archive the objects on your behalf. Add this policy in the root compartment of your tenancy. If this policy is not in place the installer will return an error. Use one statement per region.

Allow service objectstorage-<region_identifier> to manage object-family in compartment <compartment_name>

Refer to [the Object Storage documentation](#) for more information

Based on the parameters specified, the installer creates a configuration file storing the parameters that are used by RMAN. The table below shows the parameter names and their usage.

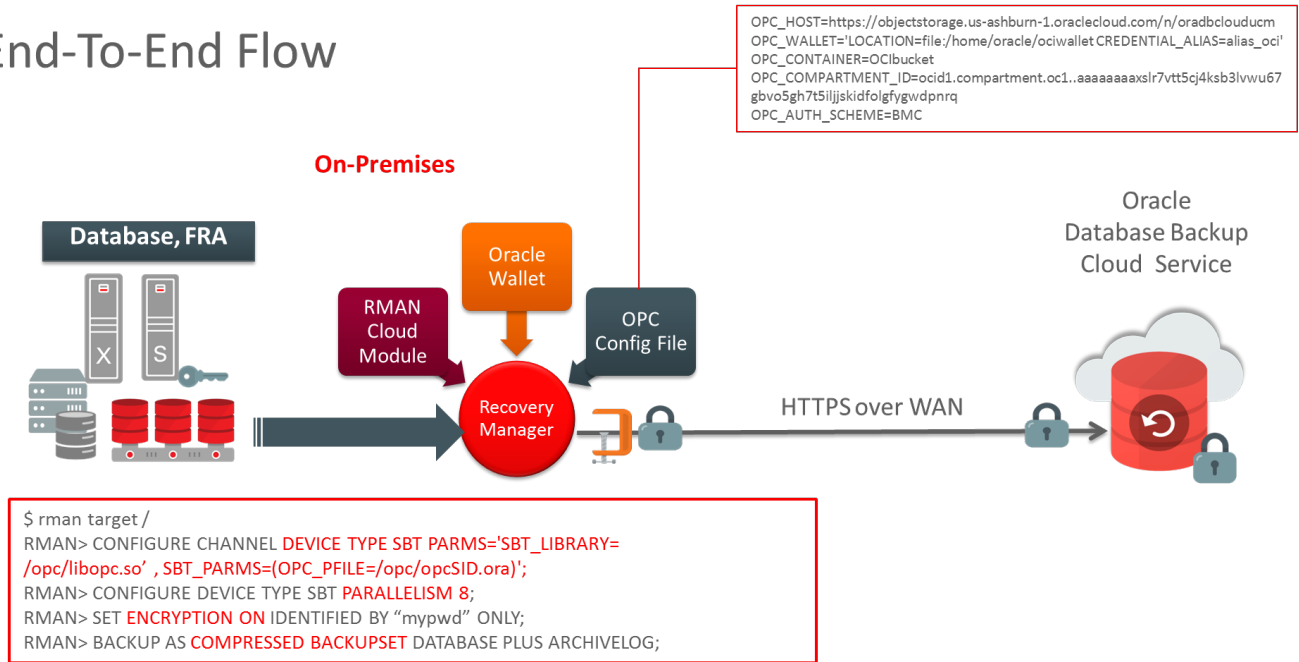
Content of the opc<sid>.ora configuration file created by the installer*:

PARAMETER NAME	DESCRIPTION
OPC_HOST	REST destination endpoint URL and storage namespace (added automatically by the installer) Ex: <a href="https://objectstorage.<region>.oraclecloud.com/n/<tenant>">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_COMPARTMENT_ID	Target compartment OCID
_OPC_TRACE_LEVEL	For debugging purposes only, set it to 100 for a complete trace in sbtio.log
OPC_AUTH_SCHEME=BMC	This is always present when the OCI native module is used

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

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

End-To-End Flow



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 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/19.1/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;
```

Both password based (like in the example above) and TDE encryption methods are supported. 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 SECTION SIZE 64G 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 SECTION SIZE 64G 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/19.1/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_mbttmlq_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](#)

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 ;
```

VALIDATION BEST PRACTICES 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 on backups
- If public network throughput is not sufficient, choose Oracle Fast Connect. Refer to cloud.oracle.com/networking
- Choose 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.

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.

APPENDIXES

USING OBJECT STORAGE REPLICATION

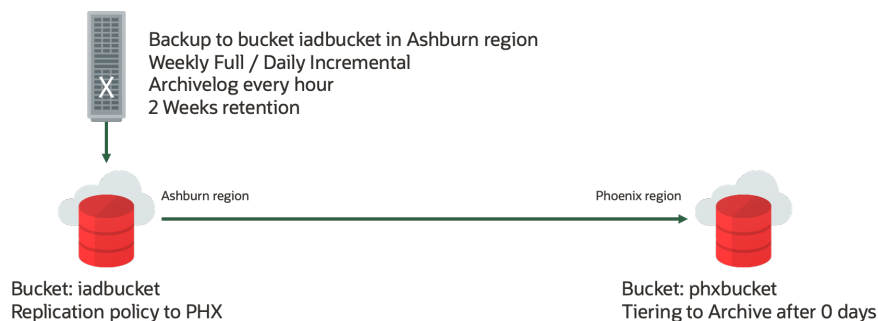
Object Storage replication can be enabled to create a copy of all the backup objects in a bucket located in the same region or in a different region.

Object Storage replication is configured by creating a replication policy on the source bucket. The replication target can be in a different region, but it must be in the same tenancy as the source bucket. A replication policy can be set up using the console or the OCI CLI following the [Object Storage documentation](#). When replication is used, restoring from the replica bucket is as easy as changing the OPC_HOST and OPC_CONTAINER parameters in the DB Backup Cloud Module configuration file. If the replica bucket requires different permissions, then the `oci_installer` must be re-run with the correct parameters in order to update the credentials wallet.

There are some considerations to keep in mind when using Object Storage replication for RMAN backups

1. When a replication policy is in place the replica bucket is read-only
2. Objects in the replica bucket are updated (created, updated, deleted) asynchronously based on source object updates.
3. The DB Cloud Backup Module only supports standard buckets, both as source or replica. Object Lifecycle Policies can be created both at the source or replica bucket to enable tiering to archive storage.
4. A deletion policy can be configured on the source bucket to automatically delete objects after a pre-defined period of time and that deletion will be reflected on the replica bucket, but a deletion policy cannot be configured on the target bucket as it's read only.

As an example, let's assume the scenario illustrated in the diagram below:



We want to send backups to a bucket in the ashburn-1 region called iadbucket and retain them for 2 weeks. For DR purposes we also want those backups to be replicated to a bucket named phxbucket in the phoenix-1 region.

The authorization needed to access the two buckets can be given to the same user or to different users. To set up this scenario we assume the two buckets are accessible to two different users, iaduser and phxuse. To configure it follow these steps:

1. Prepare the environment setting the correct authorization policies, key pairs and collecting the OIDs as explained above.
2. Create two directories to store the credential wallets, for example `~/iadwallet` and `~/phxwallet`

3. Run the oci_install.jar installer configuring access to the bucket phxbucket in Phoenix

- `java -jar oci_install.jar \`
- `-host https://objectstorage.us-phoenix-1.oraclecloud.com \`
- `-pvtKeyFile <path to phxuser private key> \`
- `-pubFingerPrint <public key fingerprint for phxuser> \`
- `-tOCID <tenacy OCID> \`
- `-cOCID <compartment OCID> \`
- `-uOCID <phxuser OCID> \`
- `-walletDir ~/phxwallet`
- `-libDir <dir where libopc.so will be downloaded to>`
- `-configfile ~/phxbucket.cfg \`
- `-bucket phxbucket`
- `-enableArchive TRUE`
- `-archiveAfterBackup "6:MONTHS"`
- `-retainAfterRestore "48:HOURS"`

4. Run the oci_install.jar installer again configuring access to the bucket iadbucket in Ashburn and specifying different wallet location and configuration file

- `java -jar oci_install.jar \`
- `-host https://objectstorage.us-ashburn-1.oraclecloud.com \`
- `-pvtKeyFile <path to iaduser private key> \`
- `-pubFingerPrint <public key fingerprint for iaduser> \`
- `-tOCID <tenacy OCID> \`
- `-cOCID <compartment OCID> \`
- `-uOCID <iaduser OCID> \`
- `-walletDir ~/iadwallet`
- `-libDir <dir where libopc.so will be downloaded to>`
- `-configfile ~/iadbucket.cfg \`

- `-bucket iadbucket`
5. Use the OCI Console to create a Replication Policy from iadbucket to phxbucket
 6. Configure the RMAN default channel allocation specifying the iadbucket.cfg as OPC_PFILE for backup
 7. To restore from phxbucket just change the channel allocation in RMAN using phxbucket.cfg as OPC_PFILE.

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 (on older tenancies)
export RCLONE_CONFIG_OCIC_AUTH=https://uscom-east-1.storage.oraclecloud.com/auth/v1.0 (on
newer tenancies, replace uscom-east-1 with the appropriate region)
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.

REFERENCES

[MAA Oracle Cloud Infrastructure Exadata Backup & Restore Best Practices using Cloud Object Storage Whitepaper](#)

[DB Backup Cloud Service on cloud.oracle.com](#)



CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com.
Outside North America, find your local office at oracle.com/contact.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2021, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Technical Brief Title

May, 2021

Author: [OPTIONAL]

Contributing Authors: [OPTIONAL]

