In this Document

Purpose

WANDISCO provides the mechanism to replicates both on-prem and cloud based big data sources into Oracle Storage Cloud. This document shows show how to integrate object storage data buckets managed with WANDISCO FUSION with LiveS3 as staging area for ADWC database tables.

Scope

This information is provided for the system administrator or database administrator to configure the database underlying storage for a set of tables. This document assumes that the ADWC instance has been provisioned based on the documentation here.

Details

The following are the steps to configure the underlying storage for an ADWC database table using Oracle Object Storage.

1) Provision a set of Compute servers in Oracle Cloud Infrastructure.

•	Once you have an Oracle Cloud account, you create 2 or more OCI Compute servers choosing the shape – which determines the CPU count and memory configuration.

•	Servers with 1-2 CPUs and 16GB is sufficient for a minimal WANDISCO FUSION installation in a non-production configuration.

•   For production use, installation of WANDISCO FUSION on more powerful servers with say 4-8 CPUs and 64 GB of memory is recommended for purposes of both performance and availability.

2) Install the WANDISCO FUSION and LiveS3 proxy server application software on the Computer servers by following the procedure outlined in this document: "WANDISCO FUSION for Object Storage in OCI"

3) Create an External Table with Object Storage Data repository using the DBMS_CLOUD PL/SQL package

Use the DBMS_CLOUD package in Oracle Autonomous Data Warehouse to load, query, and save Object Storage data. The PL/SQL package DBMS_CLOUD provides support for loading data from files in the Cloud to your tables in Autonomous Data Warehouse. This package supports loading from files in the following cloud services: Oracle Cloud Infrastructure Object Storage, Oracle Cloud Infrastructure Object Storage Classic, Azure Blob Storage, and Amazon S3.

The list below shows the DBMS_CLOUD subprograms provided with Autonomous Data Warehouse:

- COPY_DATA Procedure
- CREATE_CREDENTIAL Procedure
- CREATE_EXTERNAL_TABLE Procedure
- DELETE_FILE Procedure
- DROP_CREDENTIAL Procedure
- LIST_FILES Function
- PUT_OBJECT Procedure
- VALIDATE_EXTERNAL_TABLE Procedure

To use object storage as the underly storage for the table, we will use the CREATE_EXTERNAL_TABLE function. The first step is to establish a set of credentials to provide authentication. And the second step is to use the credentials to read and write data.

1.   Store your object store credentials using the procedure DBMS_CLOUD.CREATE_CREDENTIAL. For example:

```
SET DEFINE OFF

BEGIN

  DBMS_CLOUD.CREATE_CREDENTIAL(

    credential_name => 'OBJECT_CREDENTIALS',

    username
=>'ocid1.credential.oc1.aaaaggtx7kmqpnhuwg…yqdiomoas4vq',

    password => '1RHXlxxwBhbPDAbf9P676E5UgZyaDFRqUn86K8Je2NI='  );

END;

/
```

Note: Use the S3 object storage secret key and access key created in the OCI platform.

2.   Create an external table on top of your source files using the
procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE.`

```
BEGIN
 DBMS_CLOUD.COPY_DATA(

    table_name =>'SALES',

    credential_name =>'OBJECT_CREDENTIALS',

    file_uri_list =>
'https://myproxy.wandisco.com:8081/databucket/ADWC1/SALES/*.dat,

    format => json_object('delimiter' value ',')

    column_list =>'

    prod_id              NUMBER          NOT NULL,

    cust_id              NUMBER          NOT NULL,

    time_id              DATE            NOT NULL,

    quantity_sold        NUMBER(10,2)    NOT NULL,

    amount_sold          NUMBER(10,2)    NOT NULL'

 );
```

WANDISCO FUSION for Oracle Autonomous Data Warehouse

```
END;

/
```

Note: In this example, a wild card * is used to specify that all files in the databucket with prefix ADWC1/SALES ending in ".dat" will be used.