An Oracle White Paper
November 2013

Data Movement and the Oracle Database Cloud Service

Table of Contents

Introduction to data loading 3
Data loading options 4
  Application Express 4
  SQL Developer 4
  RESTful Web Services 5
  Which approach is right for you? 5
Data loading with SQL Developer 5
  Initialization 6
    Changing the SFTP user password 6
    Create a Cloud Connection 7
  Deployment 7
    Create a cart 8
    Monitor a deployment 9
    Reusing a data load job 10
Unloading data to another Oracle database 10

Introduction to data loading

The Oracle Database Cloud is a platform for building data-centric applications universally available in the Cloud. These applications allow users to add, change and manipulate data to extract information to allow for better operations and strategic decisions.

Some applications will begin with a clean slate – with data structures defined, but without any data in the structures. In this case, all data would be added through user interaction.

A more common scenario is an application where some amount of data already exists, although not in an Oracle Database Cloud Service. This scenario will require data to be loaded into the Oracle Database Cloud Service.

Loading data into your Database Cloud Service requires a different approach than loading data into an on-premise Oracle Database. This white paper will provide an overview of the basic methods which can be used to load data into an Oracle Database Cloud Service, as well as provide an in-depth tutorial on using one of these methods. In addition, the paper will describe how to unload data from your Database Cloud Service.

Data loading options

You can load data into your Oracle Database Cloud Service with two different tools – Oracle Application Express or SQL Developer. You can also use RESTful Web Services to create your own custom data loading process, although this option will require significantly more work.

Application Express

Application Express is a rapid application development tool included as part of the Oracle Database Cloud. Application Express runs as part of the Oracle Database Cloud and is accessed through a standard Web browser. Oracle Application Express includes SQL Workshop, a set of tools and utilities for working with data structures and data in an underlying database.

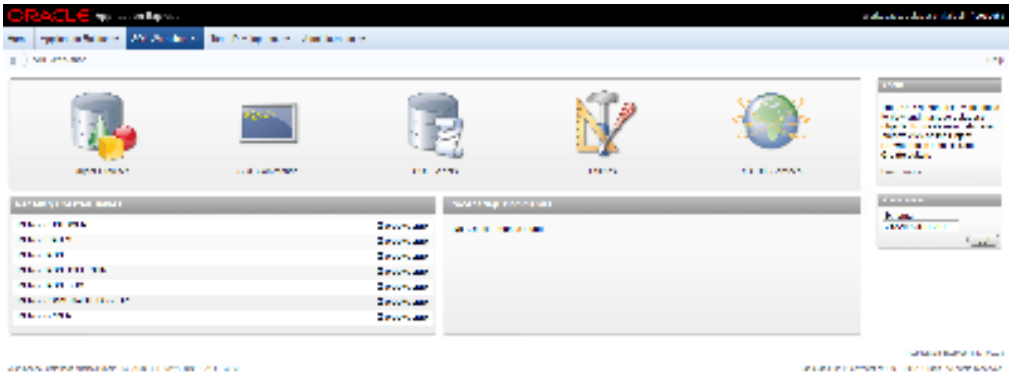

SQL Workshop contains an area which allows you to run SQL scripts. These scripts can include both SQL statements and embedded data, and can be used to load data into the Oracle Database Cloud Service. In addition, Application Express has a set of utilities that allow you to load data from a variety of formats, including CSV files, and XML files.

You can learn more about the capabilities and use of these Application Express options in the standard documentation for Application Express. You can find this documentation at this URL - http://www.oracle.com/technetwork/developer-tools/apex/application-express/apex-094287.html)

SQL Developer

SQL Developer is a popular tool that runs on a client machine, available as a no-cost download from the Oracle Technology Network. SQL Developer, versions 3.2.10 and up, can access an Oracle Database Cloud Service and allows you to access and load data and data structures into your Oracle Database Cloud Service.

SQL Developer provides a rich array of functionality, beyond the scope of this white paper, which will focus on using SQL Developer to load data into your Oracle Database Cloud Service.

RESTful Web Services

Both of the options listed above are automated utilities, which perform many steps for you within an established framework. You can also use RESTful Web Services to put data into your Oracle Database Cloud Service from outside of the Database Cloud. However, be aware that you will have to write the PL/SQL code to add the data, but also handle all other facets of the data load operation, such as marshalling data for upload, so using this method will require a fair amount of additional effort. You should always seriously consider using one of the utilities that come with your Database Cloud Service before looking at using your own custom approach with RESTful Web Service.

Which approach is right for you?

Functionally, both the SQL Workshop component of Application Express and SQL Developer give you the ability to load data flexibly and productively into your Oracle Database Cloud Service. There are some differences in the process flow for each – such as SQL Workshop running from within your browser, and SQL Developer running from a separate client program – so your particular working style or the technical limitations of your browser may be better suited to one approach or the other.

In particular, SQL Developer compresses data to be loaded before transferring it to the Oracle Database Cloud, so loading large amounts of data may be more efficient with SQL Developer. Web browsers are not designed for uploading large amounts of data over the Internet.

In addition, SQL Developer gives you the option to run any SQL script before or after the data load operation, which you can use to add functionality specific to your particular data loads.

Finally, all data loading tasks are not the same. You may choose to use one method to load your initial data, and another for incremental updates.

Data loading with SQL Developer

Data loading with SQL Developer includes a number of steps

- Choosing the tables, data and other database objects to be uploaded
- Automatic creation of a compressed file or deployment cart for transfer to the Oracle Database Cloud
- Movement of the file to a Secure FTP site
- Decompression of the file, virus scanning and

- Loading the tables, objects and data into a specific Oracle Database Cloud Service

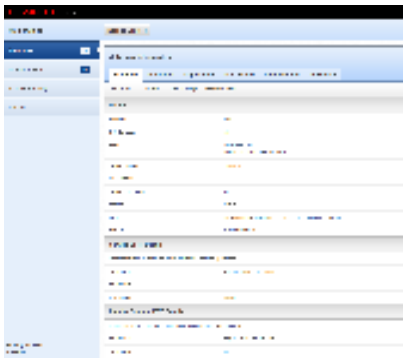

In order to perform these steps, the data load process uses two different users – one user defined for the Secure FTP (SFTP) site, and another defined to load the data in the Database Cloud Service. You can use any user with who is a Service Administrator or Database Developer to load data into your Database Cloud Service. The SFTP user for your Database Cloud Service is identified in the My Services portal[1] for the Database Cloud Service in the section label Service SFTP Details.

Using the data loading capabilities of SQL Developer has two basic phases – the initialization phase, where you configure these two users for all data load processes for a particular Database Cloud Service, and deployment, where you create "carts", which specify the tables, objects and data to be moved, and deploy these to your Database Cloud Service. The first step is done once for a Database Cloud Service, while the second step can be done one or more times, depending on your data loading needs.

Initialization

You have to perform two tasks to initially set up your Database Cloud Service to allow for data uploads from SQL Developer –

- Change the password for the SFTP user
- Create a connection to your Database Cloud Service in SQL Developer

Changing the SFTP user password

In the first initialization step for data loading, you change the password for the SFTP user defined for your Database Cloud Service and create a user to load the data into your Database Cloud Service.

Your first step is to change the password for the SFTP user listed in the My Services portal, as described above.

To change the password for this user, you must log into the Oracle Cloud Identity Management Console, which you can reach by clicking on the Identity Console button in the My Services page for the Database Cloud Service. Once you log into the Management Console, click on the Manage Users choice in the left hand box, and then click

[1] You can get to the My Service portal page for your Database Cloud Service by logging into your account at cloud.oracle.com, selecting My Services, and then choosing the link for your Database Cloud Service.

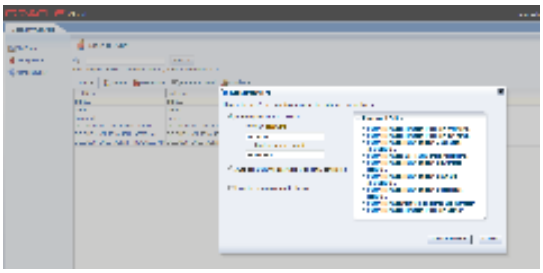

on the Search button to bring up a list of users. You should see the user identified at the Secure FTP User as one of the users.

Highlight the appropriate user and click on the Reset Password option at the toolbar on the top of the table. This action will bring up the Reset Password dialog, where you can choose to manually reset the password. Set a password for the user and click on Reset Password.

Create a Cloud Connection

The second initialization step is creating a Cloud Connection to your Database Cloud Service in SQL Developer. A Cloud Connection is slightly different from a normal connection in SQL Developer, since this type of connection uses RESTful Web Services to access your Database Cloud Service and requires slightly different information to properly access your Service.

When you start SQL Developer, you will see two main entries under Connections in the left hand browser window – Auto-Generated Local Connections and Cloud Connections. Click on the Cloud Connections entry with the right mouse button to begin the process of creating a connection to your Database Cloud Service.

This action will bring up a dialog box that prompts you for some basic information about your Database Cloud Service – a descriptive name for the Cloud Connection, the username for the database administrative user, the URL for the Service, as well as the username for the SFTP user, the hostname for the SFTP server, which was in the email you received when your Database Cloud Service was allocated, and the port for the SFTP server, which is 22.

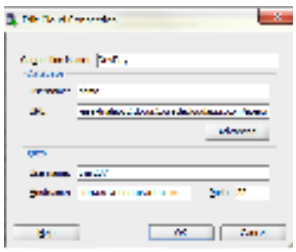

Click on the OK button once you have entered this information to create a Cloud Connection. After the Cloud Connection has been properly defined, you can open the connection to see information about the data structures in your Database Cloud Service – similar to the information you can see from a normal database connection, with one additional entry at the bottom labeled Deployments, which will track the data loading jobs you will create in the next section.

Deployment

Your Database Cloud Service is now ready to accept data loaded through SQL Developer. You will accomplish this task by using a cart, which acts as a container to hold the definitions of the content that will be loaded to your Database Cloud Service.

Create a cart

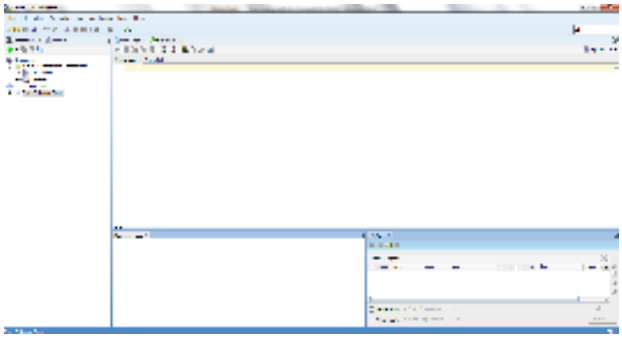

The data load process begins by creating a cart. In SQL Developer, you begin this process by selecting the Cart option from the View choice in the menu bar at the top, which will open up a cart window in the lower right hand side of SQL Developer, as shown here.

You add contents to a cart by simply dragging and dropping tables from a database. Once you have added a table to a cart, you can designate whether you only want to create the table structure in your Database Cloud Service or to create the table structure and load the data. You can also add a WHERE clause to limit the data which is uploaded to your Database Cloud Service.

You can add multiple tables to an individual cart – all the tables will be part of the same data upload deployment.

You can also add SQL scripts that will run before the data loading job starts, and after the job completes successfully, by specifying those script locations in the lower part of the cart.

Once you have added all the tables you require into the cart, you run the data load by clicking on the Deploy Cloud icon, which is on the left of the icon row and looks like a little cloud. When you click on this icon, you will be prompted to name the deployment, which cannot be more than 15 characters and can only include alphanumeric characters, underscore and dash, provide the server name, port, which should be 22, and username (filled by default) and password for the SFTP user, and a location for the compressed file that will be created as part of the overall data load job.

You can change some of the options that control the creation of the compressed file by clicking on the Deploy icon, which is to the right of the Deploy icon.

You can also specify how conditions on how the data structures and data will be deployed into the Database Cloud Service.

You also have the option of adding a SQL script to be run before the data load begins or after the data load completes. This capability allows you to include customized functionality to your load job, including extensive data validation and manipulation.

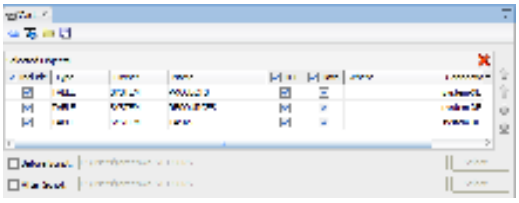

Once you have designated this information for the deployment, click on the Apply button to start your data load deployment.

Monitor a deployment

When you have clicked on the Apply button, you have started the process of data loading with SQL Developer. The first step in the process is to compress the DDL statements and data into a .ZIP file.

Once the compressed file has been created, SQL Developer transfers the file to the Secure FTP site. Once the file arrives at the SFTP site, the deployment job is listed in the Deployments section under the Cloud Connection which represents the destination for the deployment. This section lists all deployments to this Database Cloud Service. You can see more details on the progress of the deployment by clicking on the name of the deployment, which will bring up information on the status of the job in the main window of SQL Developer, as shown here.

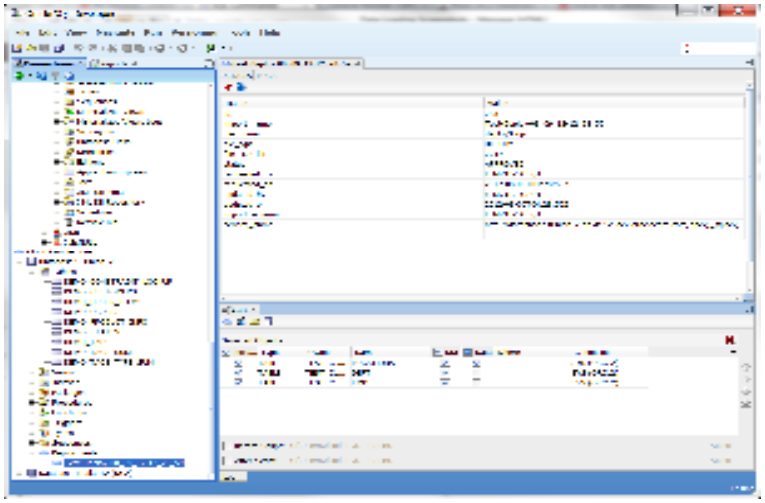

When the file arrives at the SFTP site, the import has to be approved, which is done automatically. Once the import has been approved, the status listed in the window changes to APPROVED. Once a file has been approved, a background process creates a job to run the SQL commands in the file. The background process periodically checks for incoming files, so the file may wait a few minutes before the job is created.

Once the job is created, the data load starts. Once the data load is completed, the status changes to either PROCESSED or ERRORED, depending on the outcome. (You can refresh the status with the Refresh button at the top of the panel.) If you receive a status of FAILED VERIFICATION, there was probably a problem with the user name and password you used for the SFTP user.

Once your data loading job reaches the status of PROCESSING, you can track the progress of the job by simply refreshing the list of tables in the designated Cloud Connection. Once a table is created, it will appear as part of the Database Cloud Service list of tables. You can also use the Data Browser in the SQL Workshop area of Application Express to both see when a table has been created and also track the progress of the load by clicking on the Data tab and then the Count Rows button to see how many rows are currently in the table.

Keep in mind that the import job only commits new row insertions periodically, so the number of rows in a table will advance in line with the inserted and committed rows.

After a data load job has completed its run, you can click on the Logs tab to see a listing of the completed steps. The load process also puts log files into the SFTP site. You can connect to the site using a standard tool, such as the free FileZilla, and look in the download directory. You will see a .ZIP file with the prefix of LOG_ and then the name of your deployment, as well as log files for each individual component loaded. If there are rows which are not inserted as part of the load, the rows will be noted in the overall log file as well as included in a separate file which will contain the word 'bad'.

Reusing a data load job

You can save the contents of a cart by clicking on the Save icon in the menu bar of the Cart window. By saving a cart, you can re-run the deployment by loading the cart and following the steps described above.

You also have the option of restarting a deployment by right clicking on the deployment name and selecting the Restart option. You would restart a deployment if the deployment failed to properly load table structures or data for some correctable reason, such as duplicate values for a unique primary key. Restarting a deployment reruns the entire job, so you would have to manually clean up any partial results in your Database Cloud Service before restarting the deployment.

Unloading data to another Oracle database

The Database Cloud Service makes it easy to unload data from your Service. You start from cloud.oracle.com and then go to the Cloud Management Console by clicking on My Services, logging in as a Service Administrator. You will see all of your Services listed, and you can click on the name of any of your Database Services to go to a detail page for the Service.

On this page, one of the tabs is for Data Export, as shown in this figure. To begin a data export, simply click on the Export Data button at the top of the table listing. This action will take you to a page where you can indicate whether you want to include the data in your data export or simply export SQL statements to recreate the data structures from your Database Cloud Service.

After you have made your selections, click on the Begin Data Export button, which will return you to the listing of data exports. Your data export is given a name and an initial status of Requested. A background process periodically scans for data export requests and creates a file with the SQL statements for the request. Once the export job has begun, the status changes to Processing. When the export file has been created, the status changes to Processed.

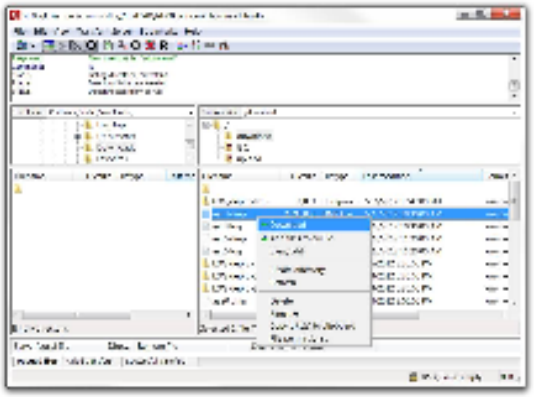

Once the export file has been created, it is available on the Secure FTP server identified for your Database Cloud Service. You can access the file with any standard FTP tool, such as FileZilla, a free tool shown here.

The export file is an Oracle Data Pump formatted dump file with an extension of .dmp. Data Pump is a standard Oracle utility. You can use Data Pump to create data structures in any Oracle database and load data, where the data is present

in the export file.

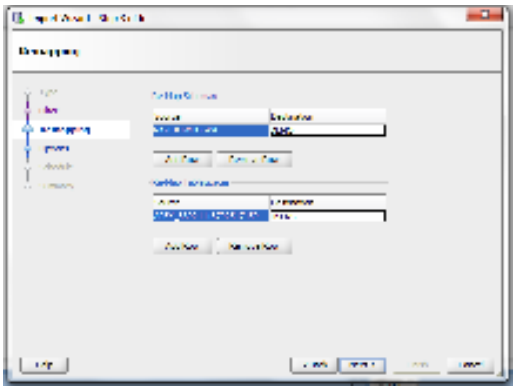

You can also use the Data Pump Import Wizard which is a part of SQL Developer. To access this Wizard, you have to open a DBA window and connect to the non-Database Cloud system to which you want to load the data. In the DBA window, you can open the section labeled Data Pump and right click on the Import Jobs entry. One of the choices will be Data Pump Import Wizard, which will start the import process.

In the first page, you can give your import job a name, and select whether you want to import the entire schema or just some tables. The Change Import Files area lets you specify where the .dmp file created by the export process resides. You have a limited number of choices for this location, based on logical identifiers like DATA_PUMP_DIR[2], a directory where Data Pump files are normally stored. Select one of the logical locations and then double click on the File Names column to enter the appropriate file name.

Once you have entered the file name for the .dmp file, you can click on Next, which will allow you to select tables or schemas to import, depending on the type of import you specified in the previous page.

The next page requires some important information. When you export data from your Database Cloud Service, the data, by default, is loaded into a schema with the same name as the schema for your Database Cloud Service, and into a tablespace which exists in the Database Cloud. You will probably want to map the incoming data to a different schema and tablespace. You can do this on this Remapping page, where you can add a row to remap the schema name and the tablespace name. As with the earlier page, you select the schema name or tablespace name in the left hand column and double click to enter the name of the destinations.

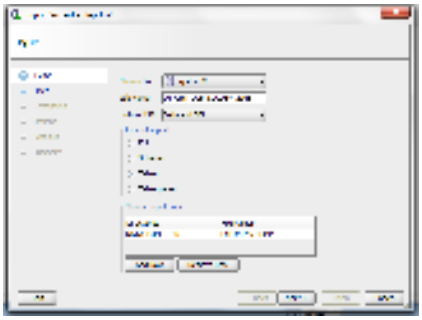

After clicking Next in the Remapping page, you can designate a directory for the log files from the Data Pump import job and other options. On the Next page, you can indicate a name for the job and if it should execute immediately. Once you move to the Next page, you can click Finish to run the job.

You can use this method to move the entire contents of your Database Cloud Service to another Oracle Database, and then eliminate any unwanted data in the new destination.

---

[2] You can find the pathname to the DATA_PUMP_DIR with the SQL query SELECT directory_path from dba_directories WHERE directory_name = 'DATA_PUMP_DIR'

Oracle is committed to developing practices and products that help protect the environment

Data Movement
November 2013
Author: Rick Greenwald

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Copyright © 2013, 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.

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

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0110