



1PLACE 1CLOUD

Oracle to Amazon Redshift

What challenges does 1Place meet?

Did you know that Databases typically make up only 12% of the overall technology footprint when moving from on-premise systems to Amazon, Azure and Google Clouds, but are one of the most complex and resources and time hungry deliverables?

Many organisations are also technically indebted, especially those who have invested in Mission Critical Tier 1 Applications supported by on-premise databases.

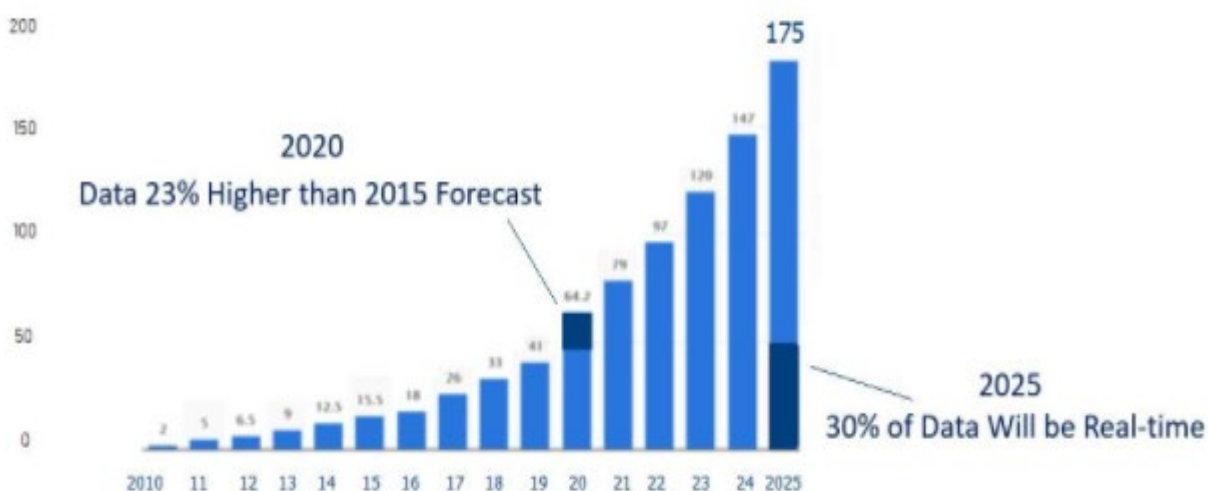
Significant complexity arises when those databases that support these applications cannot be taken offline without causing significant operational disruption either to their customers or internal users.

Legacy Database Debt and Real-Time Data

Technical database debt is also being compounded by real-time data growth. The world is moving towards greater data growth at an almost exponential rate and within this time frame real-time data is becoming a significant portion of the data captured across many organisations and those failing to utilise this data are losing business competitiveness.

Real-Time Data Growth

Volume of Data captured, created, copied and used from 2010 - 2015 (In Zetta Bytes)



2015 Year Prediction 23% Out in Just 5 Years

So why use 1Place?

Our Cloud Engineering Consultants have identified that organisations need industrial strength Cloud Migration Tools with the ability to migrate large volumes of active and real-time data held within their databases quickly and easily.

Organisations need to be able to manage their private, hybrid and multi-cloud platforms, with real-time data and not be locked into legacy databases.

Deliver Database freedom and reduce technical Debt

Organisations should be free to move away from being technically in debt with high-cost, legacy vendor licenses. 1Place enables Databases such as Oracle to convert to PostgreSQL or Hadoop that can be hosted in the cloud and managed as a unified data set at a much lower cost.

This approach includes moving Oracle large volumes of data to the Amazon RedShift data warehouse.

How can you achieve No downtime and Real-Time Migrations?

Critical legacy applications that are core to an organisation's operations cannot suffer downtime or afford data loss and moving them needs significant resources, investment, and planning. How can you make this possible with 1Place.

Summary

Organisations need to migrate larger volumes of data in real-time with zero downtime and move away from technically indebted, legacy databases and applications. There are limited approaches that are cost effective and can deliver within aggressive time frames, whilst keeping resources costs low.

1Place is built for organisations wanting to fully control their database migrations and management easily and cost effectively.

What is 1Place?

1Place supports Mission Critical Applications and Databases, whether Private, Hybrid and Multi-Cloud Database Architectures. Where data needs to be kept synchronised, managed and monitored in real-time, all of the time 1Place delivers Cloud unification.

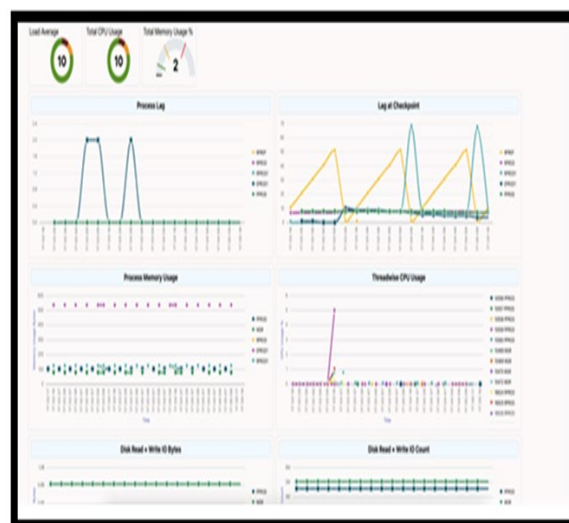
To reduce technical debt and manage real-time data 1Place can convert database sources from legacy, high license cost databases such as Oracle, to open source, low-cost databases

such as PostgreSQL and Hadoop. For large datasets 1Place can migrate to products such as Amazon Redshift.

Single View Dashboard

1PLACE empowers organisations to seamlessly manage their Cloud migrations and environments from a single dashboard and Cloud control panel all in real-time.

Rather than requiring expensive Cloud DBA's and Admins to surf between many screens to keep track of migrations and data synchronisation 1Place present all of the required data in a single View with alerts and simplified troubleshooting.



Reduced Skills, whilst retaining expertise

1Place removes the level of expertise required to operate high-value, high volume database migrations such as Oracle/MySQL/MS SQL Server/PostgreSQL and automates the installation and troubleshooting functions by up to 50%.

Reduce Resources Energy and Risk by 50%

As an example, [1PLACE Migrations are 50% faster](#), reduce risk, require much lower technical skills, enable real-time troubleshooting and save significant costs by reducing time to resolve issues and cloud storage costs.

As migrations are performed faster, risk is significantly reduced as the time taken to perform the migration is much lower reducing the time for things to go wrong.

For organisations wanting to reduce their Carbon Footprint and meet their green credential targets 1Place significantly lowers the energy required to migrate large volumes of real-time data.

How does 1Place meet today's data challenges?

1Place is a comprehensive range of cloud native functionality and tooling that provides Cloud DBa's and Support personnel with enterprise cloud migration and management technology.

1Place combines database migration, with database conversion, cloud migration and cloud management across private, hybrid and multi-cloud environments with proven technology, automation and pre-configured installation scripts and integrations with Cloud monitoring tools such as Amazons Cloud Watch.

What Can 1Place do?

SQL to No-SQL Conversions in Real-time

1Place supports Oracle to Hadoop and/or Mongo DB data conversions, in real-time allowing large data sets to be converted and managed without the cost overhead and resources required to manage SQL databases such as Oracle.

This includes the ability to take a data set contained with an Enterprise Oracle Database and load the data into Amazon Redshift.

What is Amazon Redshift?

[Amazon Redshift](#) is Amazon Redshift is a data warehouse product which forms part of the larger cloud-computing platform Amazon Web Services and is built on top of massive parallel processing (MPP) technology and can handle large scale data sets and database migrations.



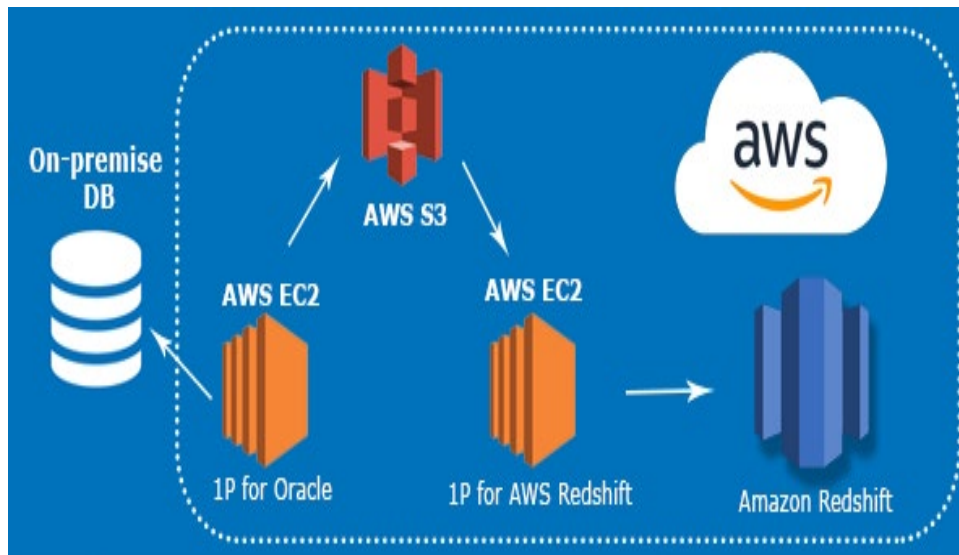
Redshift differs from Amazon's other hosted database offering, Amazon RDS, in its ability to handle analytic workloads on big data data sets stored by a column-oriented DBMS principle.

In terms of size Redshift allows up to 16 petabytes of data on a cluster compared to Amazon RDS's maximum database size of 16TB which provides you with the size and scale of Redshift and it's capabilities.

Oracle to Redshift Migration

Introduction

In this example we are performing a data load from an Oracle Pluggable Database residing On-premises to Amazon Serverless Redshift in AWS Cloud integrated with Oracle GoldenGate.



AWS Redshift table creation

By default 1Place Automatically creates the target Redshift tables by identifying the source column datatypes and maps it with Redshift Datatypes. To achieve this the schemas have to be pre-created. The following diagram describes the environment details and software versions.

Source Environment – On-premises:

- Oracle 19c PDB
- Oracle Goldengate 19c for Oracle
- 1Place Agent running on the Goldengate server

Target Environment – AWS:

- Amazon Serverless Redshift
- 1Place Agent running on the EC2 server

The target EC2 environment where 1Place for AWS Redshift runs, is created with all its pre-requisites using an Automated AWS CloudFormation template available through AWS Market Place.

Please note the EC2 instance is tuned with all the performance kernel parameters needed for Network Packet transfers.

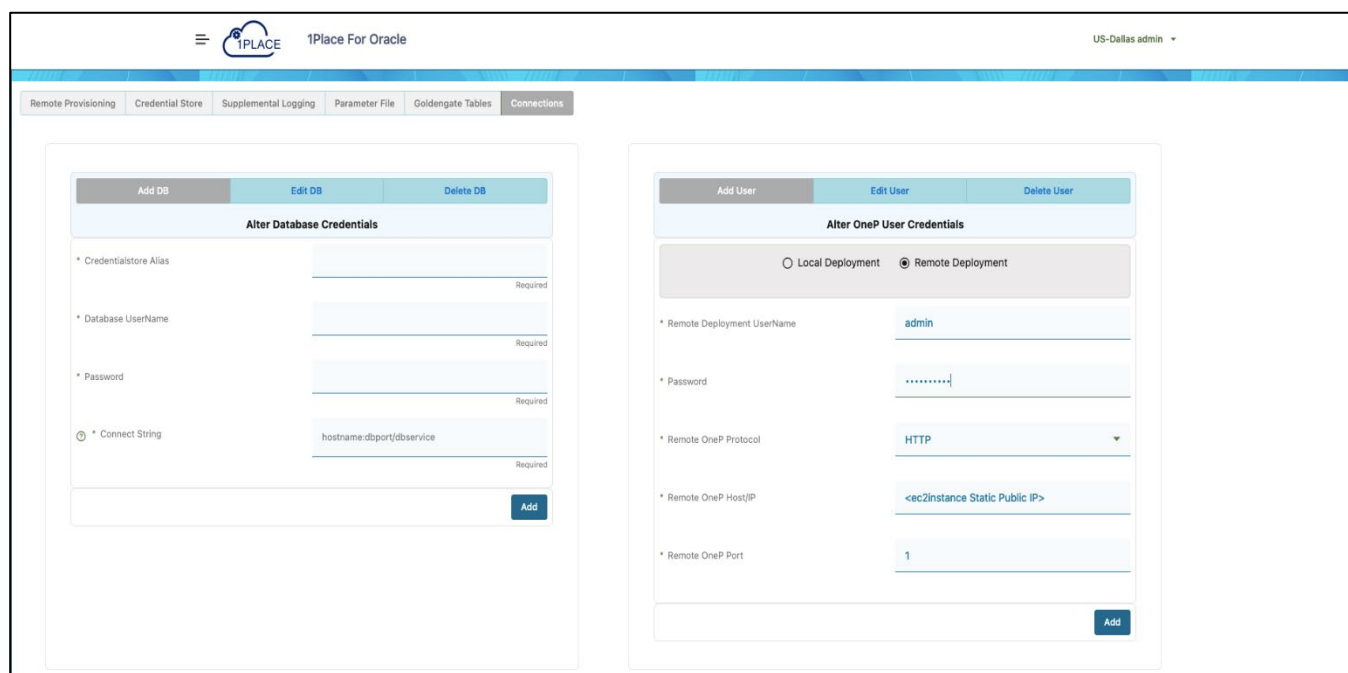
These are 2 different products residing in 2 different Datacenters/Cloud. 1Place unifies the management of these 2 products under single console which makes the management and monitoring of these 2 different environments easy.

Step -1 Initial Data load

Let's take an example of doing an initial Data load from Oracle Database to AWS Redshift and see how 1Place makes the entire process easy.

1Place uses individual Deployments to identify individual Oracle Goldengate Homes which are predefined during the 1Place setup. In this case Source Deployment is US-Dallas hosting Oracle Goldengate 19c for Oracle and 1place Agent for Oracle also Target Deployment is US-Alaska hosting 1Place Agent for Redshift.

Now login to the Source Deployment 1Place console using the credentials created during the setup and integrate the Target Deployment using the Remote Deployment option as shown below using the Connections Page.

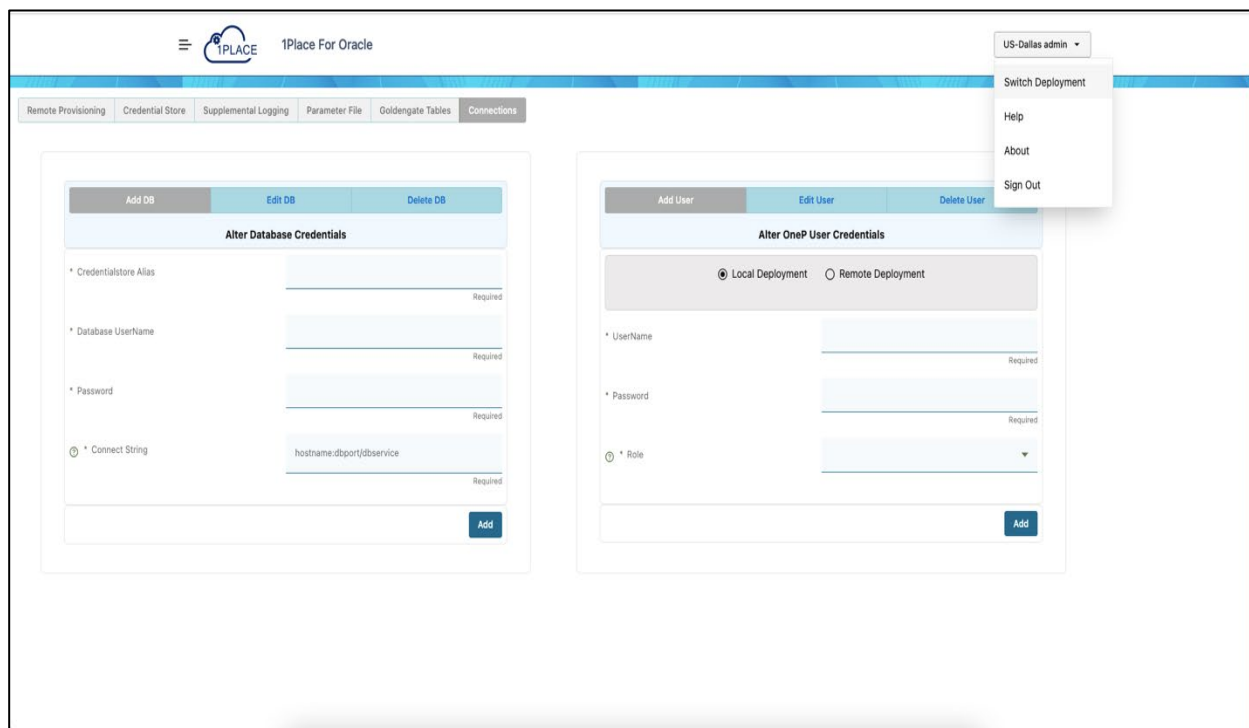


The screenshot displays the 1Place console interface for '1Place For Oracle'. The top navigation bar includes 'Remote Provisioning', 'Credential Store', 'Supplemental Logging', 'Parameter File', 'Goldengate Tables', and 'Connections'. The user is logged in as 'US-Dallas admin'. The 'Connections' page is active, showing two main panels:

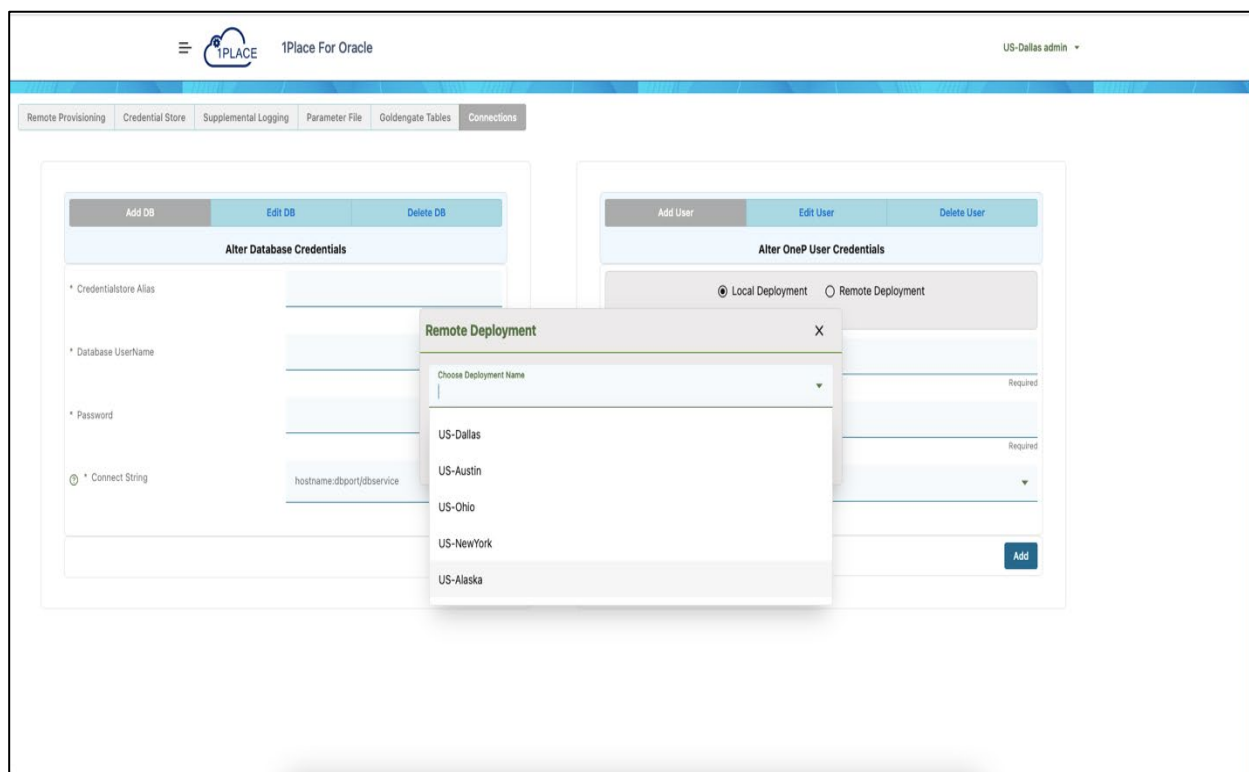
- Alter Database Credentials:** This panel has tabs for 'Add DB', 'Edit DB', and 'Delete DB'. It contains fields for 'Credentialstore Alias', 'Database UserName', 'Password', and 'Connect String' (with a default value of 'hostname:dbport/dbservice'). Each field is marked as 'Required'. An 'Add' button is at the bottom right.
- Alter OneP User Credentials:** This panel has tabs for 'Add User', 'Edit User', and 'Delete User'. It features radio buttons for 'Local Deployment' and 'Remote Deployment', with 'Remote Deployment' selected. Fields include 'Remote Deployment UserName' (value: 'admin'), 'Password', 'Remote OneP Protocol' (value: 'HTTP'), 'Remote OneP Host/IP' (value: '<ec2instance Static Public IP>'), and 'Remote OneP Port' (value: '1'). An 'Add' button is at the bottom right.

Once integrated the 2 products can be unified under a single 1Place console and managed and monitored through AWS Cloud Watch.

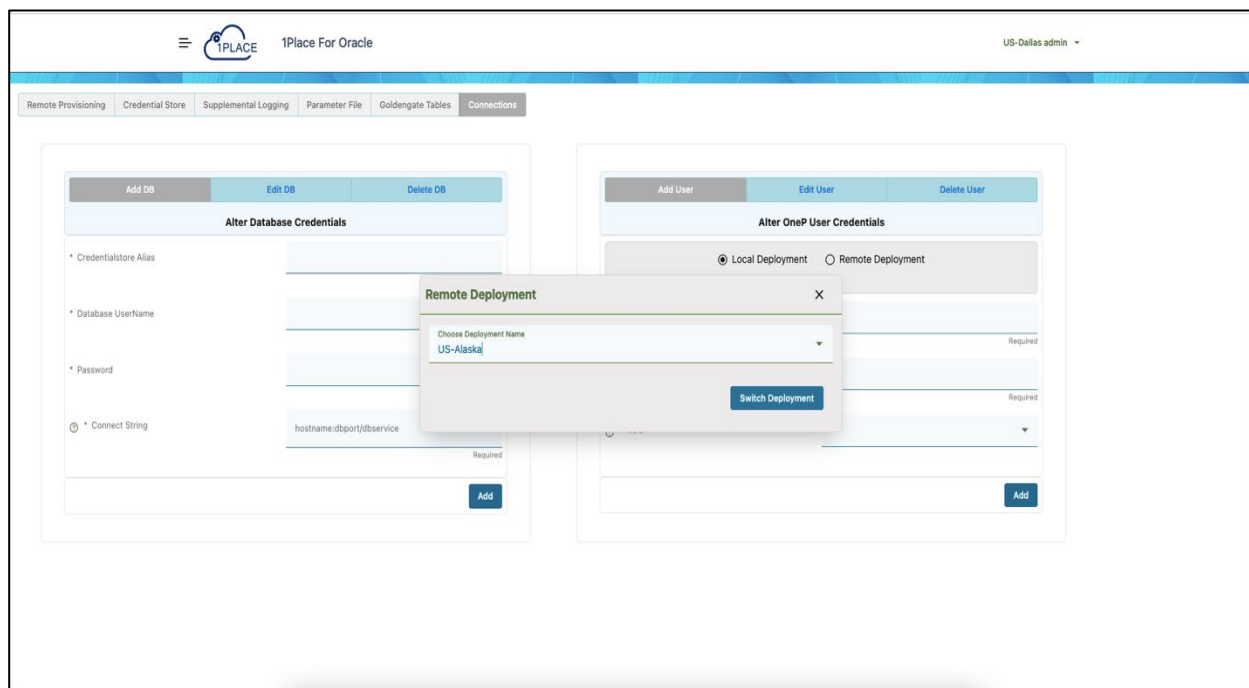
Step 2 - Go to the Deployment Name on the top right corner and click on Switch Deployment.



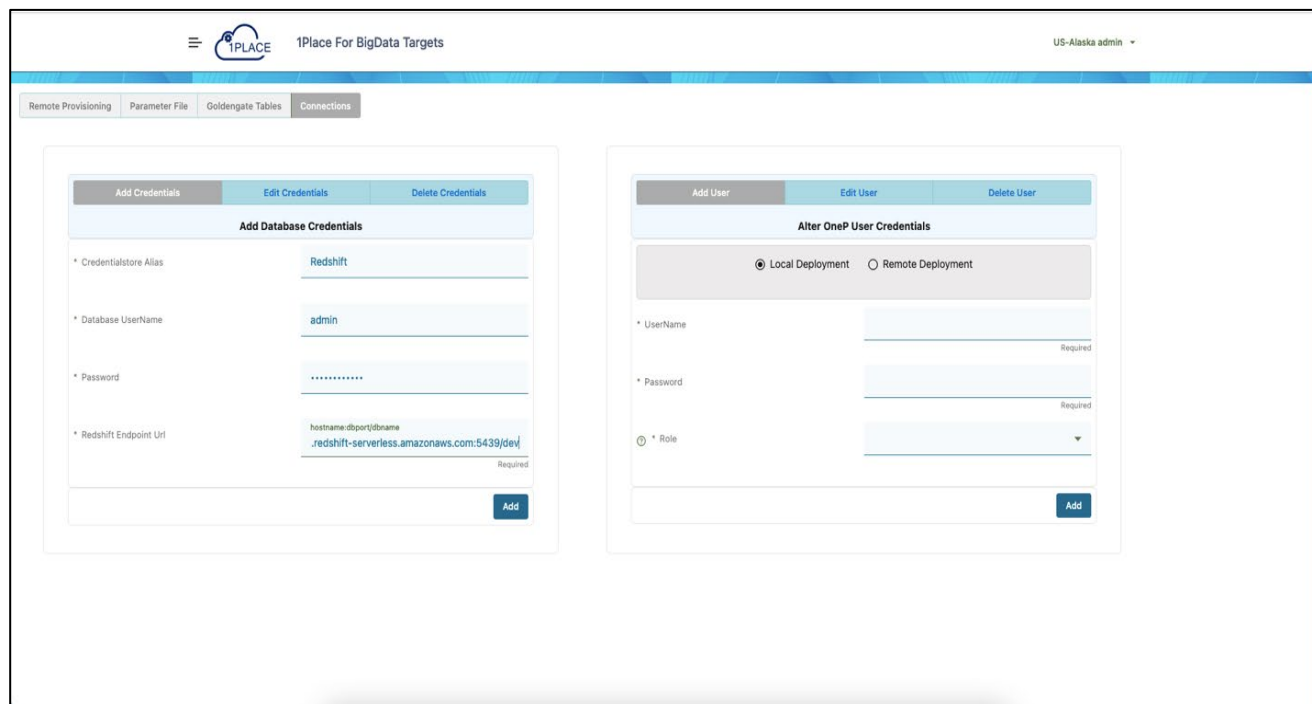
Step 3 - From the drop down choose US-Alaska which is the target Deployment.



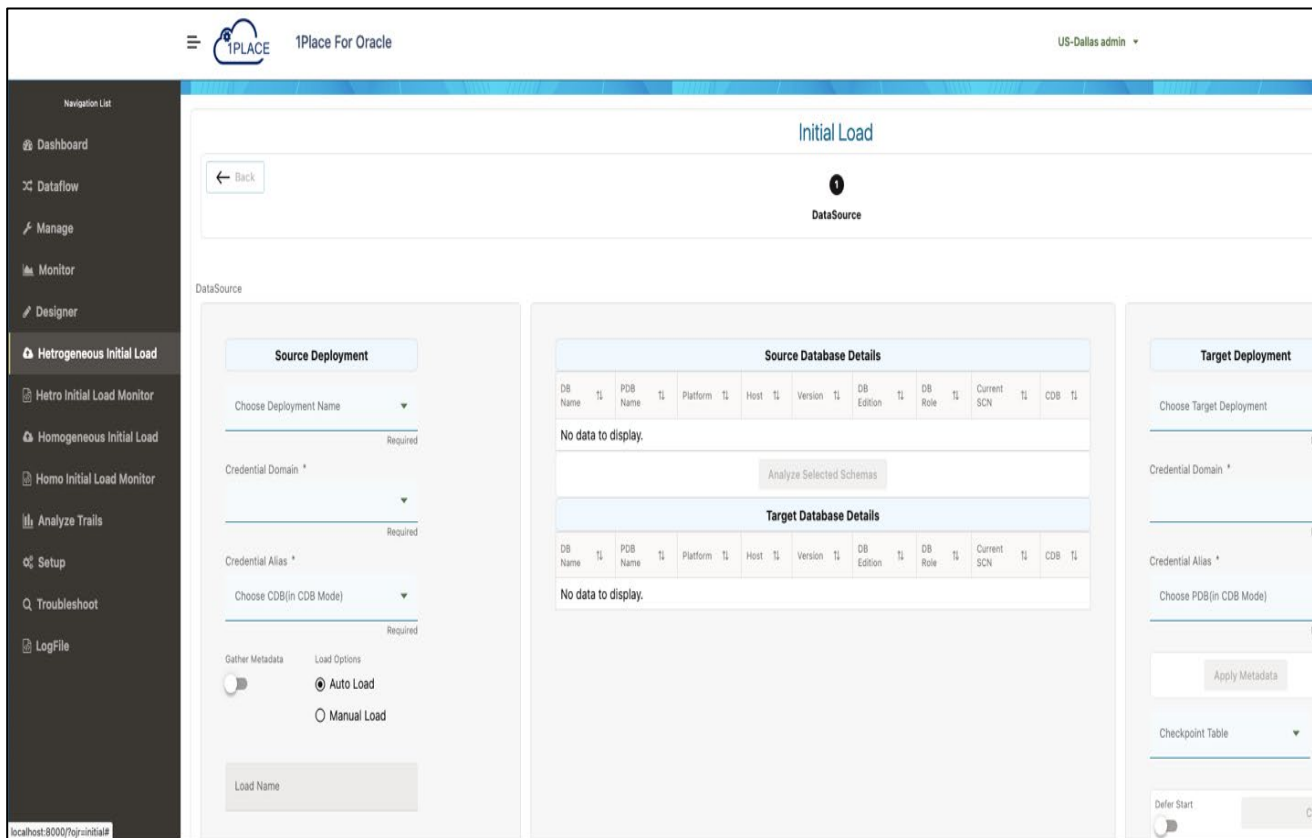
Step-4 Switch the Deployment to US-Alaska.



Step – 5 Once the deployment is switched to US-Alaska click on the connections tab and add the Redshift Database credential Alias as shown below.



Step 6 - Now go to the Heterogeneous Initial Load menu option.



Step 7 - In the Source Deployment do the below in the sequence order (Screenshot below)

1. Select the deployment as US-Dallas which Hosts the Oracle Goldengate for Oracle.
2. Select the Oracle Database Credential Alias from the 1Place console which is pre-created through the Credential store Page in 1Place. In this case since it is a Multitenant database, please choose the CDB alias.
3. Select the Pluggable database name from which you do the initial load.
4. Select the Schemas which are part of the initial load from the PDB.
5. Provide a Load Name and click on Analyze Schema button.
6. By Analyzing the tables, it will prepare the metadata of the tables which are part of the selected schema which will be used to create the tables on the Target Redshift Database. Please Note the datatype mapping is automatically done by 1Place Agent.

Initial Load Screen

← Back Next →

1
Data Source

Source Deployment

Choose Deployment Name
US-Dallas

Credential Domain *
OnPremise

Credential Alias *
Choose CDB in CDB Mode
OGG-CDB

Choose PDB Name
PDB1

Choose Schema Name
OSMDBA

Global Variables Load Options

Auto Load Manual Load

Source Database Details

DB Name	PDB Name	Platform	Host	Version	DB Edition	DB Size	Current SCN
ORCLCDB	PDB\$SEED	Linux x86_64-bit	echoapps390.com.net	19.0.0.0.0	Enterprise	PRIMARY	88236681
ORCLCDB	PDB1	Linux x86_64-bit	echoapps390.com.net	19.0.0.0.0	Enterprise	PRIMARY	88236682
ORCLCDB	PDB2	Linux x86_64-bit	echoapps390.com.net	19.0.0.0.0	Enterprise	PRIMARY	88236683

Analyze Selected Schemas

Target Database Details

DB Name	PDB Name	Platform	Host	Version	DB Edition	DB Size	Current SCN	CDB
No data to display.								

Target Deployment

Choose Target Deployment
Required

Credential Domain *
Required

Credential Alias *
Required

Choose PDB in CDB Mode
Required

Apply Defaults

Checkpoint State +

Table Split Create All

Load Name
AWR

Select the tables to exclude from initial Load

Schema Table Name	Row Count	Average Row Size	Analysis Date
<input type="checkbox"/> OSMDBA.OWB_TRANSACTION	122829373	13288	09-AUG-21
<input type="checkbox"/> OSMDBA.OWB_FINANCE_TXN	77000000	8348	09-AUG-21
<input type="checkbox"/> OSMDBA.OWB_HR_TXN	23849001	2354	09-AUG-21
<input type="checkbox"/> OSMDBA.POS_AVOID	680000	79	09-AUG-21
<input type="checkbox"/> OSMDBA.ASSD	1	0	19-AUG-21
<input type="checkbox"/> OSMDBA.OWB_PIC_DATA	0	0	09-AUG-21
<input type="checkbox"/> OSMDBA.IDENTITY_TEST_TAB	2	0	09-AUG-21
<input type="checkbox"/> OSMDBA.MANU/TEST	0	0	09-AUG-21
<input type="checkbox"/> OSMDBA.TS	0	0	16-NOV-21
<input type="checkbox"/> OSMDBA.TEST	6	0	05-NOV-21
Total Average Size		34267	

Page 1 (1-10 of at least 11 items) | (| 1 2 3 ...)

Tables to exclude from initial Load

Tables qualify for split

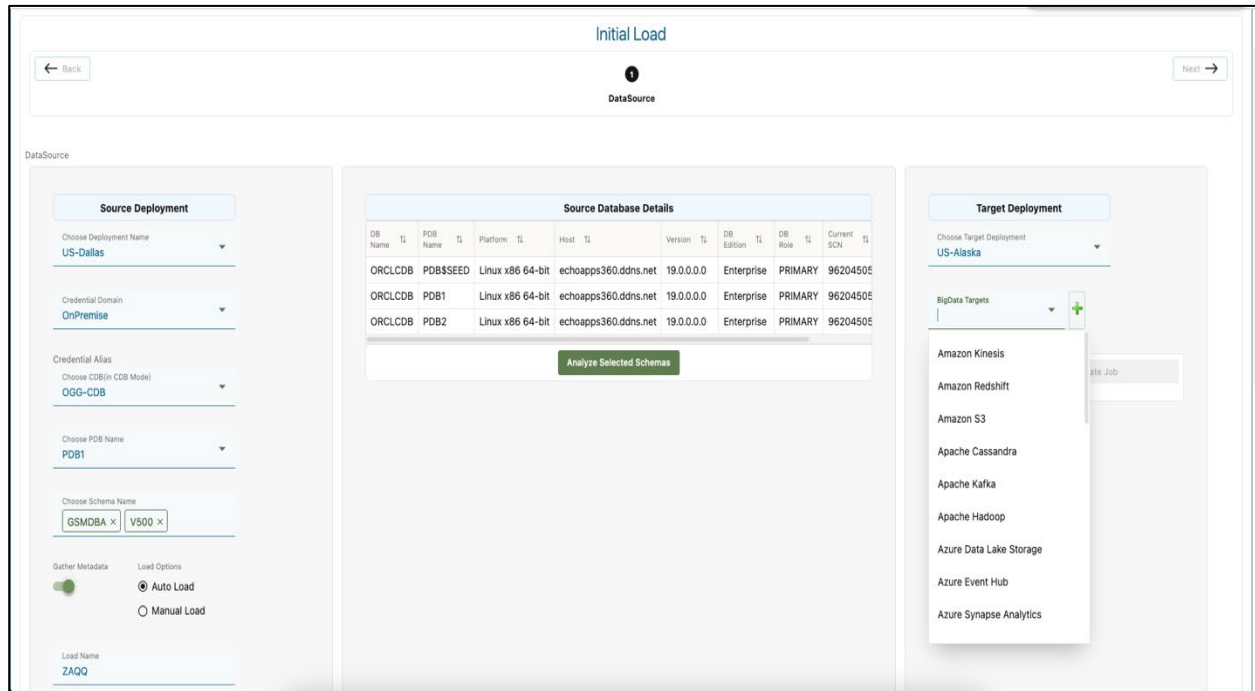
- OSMDBA.OWB_FINANCE_TXN
- OSMDBA.OWB_HR_TXN
- OSMDBA.OWB_TRANSACTION

In the above screen you can exclude any objects that don't need to be a part of the initial Load.

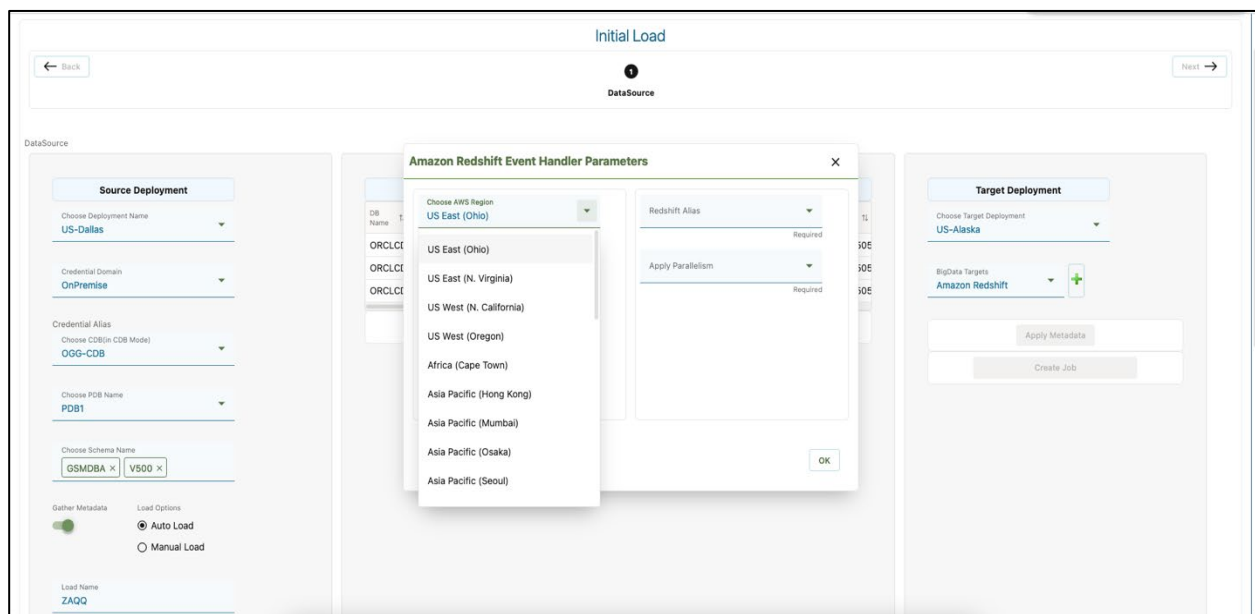
Copyright 2022@1Place | www.1place1cloud.com

Page 11 | 14

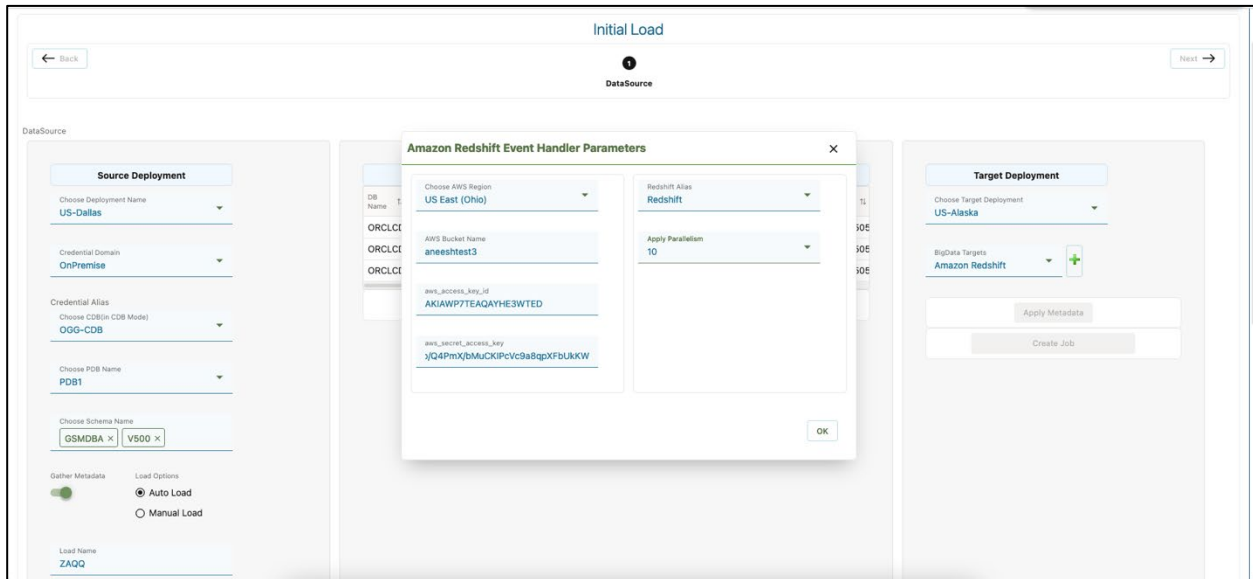
Step 8 - In the target Deployment choose the Target deployment as US-Alaska which Hosts the 1Place Agent for Amazon Redshift. And choose the Big Data target as Amazon Redshift from the dropdown list.



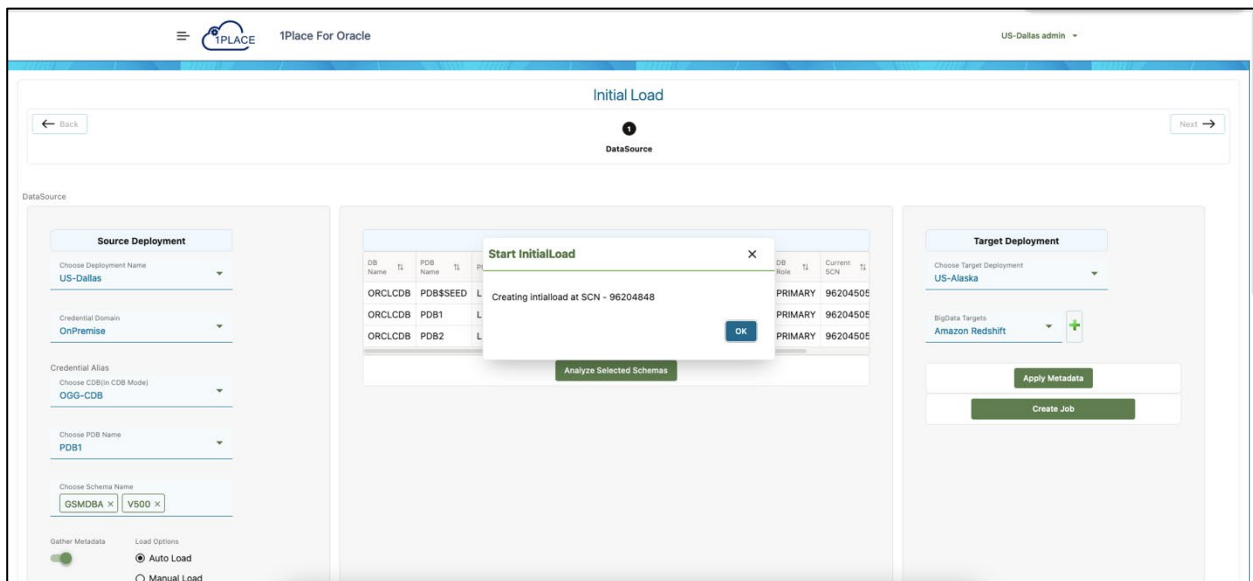
Step 9 - Click on the Add Button and it will pop up the parameters needed for transferring the extracted data in compressed format to the S3. Choose the AWS Region of the S3 bucket and Amazon Redshift. It is mandatory to have the S3 Bucket in the same region as of the Amazon Redshift. For S3 authentication and Redshift authentication need to provide the AWS Secret Key ID and AWS Secret Access Key.



Step 10 - Provide the Redshift alias which was created earlier in this document. Choose the parallelism based on the number of CPU core's available on the EC2 instance hosting the 1Place for AWS Redshift.



Step 11 - Click OK and click Create Job. This will create an Initial Load Job at a particular SCN.



Step 12 - The Initial Load Processes has 2 components.

- 1) Extract – Which extracts the source Data and upload the compressed data to S3.
- 2) Replicat – Reads the data from the S3 and loads into the AWS.

Step 13 - In the 1Place console you can monitor the number of records extracted from Oracle database and loaded to the Amazon Redshift.

Schema	Table Name	Target Rows	Extracting Count	Extract Elapsed Seconds	Extracting Rate/Sec	Replicating Count	Overall Rate/Sec	Total Elapsed Seconds	ETA
PDB1.GSMDBA	GWR_PIC_DATA	0	0	0	0	0	0		0
PDB1.GSMDBA	ASSD	1	0	0	0	0	0		0
PDB1.GSMDBA	POS_WIION	650000	0	0	0	0	0		0
PDB1.GSMDBA	GWR_TRANSACTION	122828373	4989539	56	89099	4989539	57351	1m27s	34m15s
PDB1.GSMDBA	GWR_FINANCE_TXN	77000000	0	0	0	0	0		0
PDB1.GSMDBA	GWR_HR_TXN	23649031	0	0	0	0	0		0
PDB1.GSMDBA	MANJUTEST	0	0	0	0	0	0		0
PDB1.GSMDBA	T_RAW	0	0	0	0	0	0		0
PDB1.GSMDBA	TEST	6	0	0	0	0	0		0
PDB1.GSMDBA	TE	0	0	0	0	0	0		0
PDB1.GSMDBA	IDENTITY_TEST_TAB	2	0	0	0	0	0		0
PDB1.V500	TEST_MIX1	1	0	0	0	0	0		0
PDB1.V500	FREQUENCY_SCHEDULE	0	0	0	0	0	0		0
PDB1.V500	TEST_MIX2	1	0	0	0	0	0		0
PDB1.V500	TE	0	0	0	0	0	0		0

Summary

The migration process to move Oracle to Redshift as outlined in this document shows how intuitive and straightforward the migration is using 1Place.

For further help, advice please contact the 1Place Offices on 0121 594 0686, email hello@1place1cloud.com or visit www.1place1cloud.com