Oracle® Database Using Oracle Data Transforms





Oracle Database Using Oracle Data Transforms,

F51654-20

Copyright © 2021, 2024, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Contents

Preface

Audi	ence	V
Doc	umentation Accessibility	V
Rela	ted Documents	V
Con	ventions	vi
Ov	erview	
1.1	Introduction to Oracle Data Transforms	1-1
1.2	Supported Browsers	1-2
1.3	Supported Database Functions	1-2
Pre	erequisites	
2.1	Create an SSH/RSA Key	2-1
2.2	Create Dynamic Group and Policies	2-2
De	ploy	
3.1	Select the Product	3-1
3.2	Create the Oracle Data Transforms Instance	3-2
3.3	Add Ingress Rules to the Default Security List	3-4
3.4	Provision the Oracle Data Transforms Metadata Repository	3-4
Aco	cess	
4.1	Connect to Oracle Data Transforms Instance	4-1
4	4.1.1 Connect to Oracle Data Transforms Instance from Linux	4-1
4	4.1.2 Connect to Oracle Data Transforms Instance from Windows	4-2
4.2	Access the Oracle Data Transforms Compute Instance through VNC	4-3
4.3	Access Oracle Data Transforms	4-2



5 Connections

	5.1	Work with Connections	5-1
	5.2	Supported Connection Types	5-3
	5.3	Create Custom Connectors	5-13
	5.4	Create a Data Transforms Connection for Remote Data Load	5-14
	5.5	Create an Oracle Business Intelligence Cloud Connector Connection	5-15
	5.6	Create a Delta Share Connection	5-16
	5.7	Create an Oracle Financials Cloud Connection	5-19
	5.8	Create and Use an Oracle NetSuite Connection	5-21
	5.9	Create an Oracle Object Storage Connection	5-23
	5.10	Create a REST Server Connection	5-25
6	Dat	a Loads	
	6.1	Create a Data Load	6-1
	6.2	Run a Data Load	6-3
	6.3	Monitor Status of Data Loads, Data Flows, and Workflows	6-6
7	Dat	a Entities	
	7.1	Work with Data Entities	7-1
	7.2	Import Data Entities	7-2
	7.3	Create Data Entities	7-3
	7.4	Create Data Entities within the Data Flow editor	7-5
	7.5	View Statistics of Data Entities	7-6
8	Pro	jects	
	8.1	Work with Projects	8-1
9	Dat	a Flows	
	9.1	About Data Flow Editor	9-1
	9.2	Create a Data Flow	9-3
	9.3	Add Components	9-4
	9.4	Component Properties	9-5
	9.5	Map Data Columns	9-6
	9.6	Validate and Execute a Data Flow	9-7
	9.7	Schedule Data Flows or Workflows	9-7
	9.8	Monitor Status of Data Loads, Data Flows, and Workflows	9-9



Workflows 10 10.1 10-1 Introduction to Workflows 10.2 Create a New Workflow 10-2 10.3 Schedule Data Flows or Workflows 10-4 10.4 Monitor Status of Data Loads, Data Flows, and Workflows 10-6 Jobs 11 11.1 Create and Manage Jobs 11-1 12 Users 12.1 12-1 Create Users 12-2 12.2 Change User Passwords 12.3 12-3 **Delete Users** 13 **Autonomous Database** 13.1 Manually Register the Created Oracle Data Transforms Instance to Autonomous Database 13-1 13.2 Map Autonomous Database User to Data Transforms User 13-2 13.3 Unregister the Oracle Data Transforms Instance from Autonomous Database 13-3 13.4 Access Oracle Data Transforms From Database Actions page 13-4 13.5 Access Oracle Data Transforms From Data Studio 13-5 Enable Access to Private Data Sources from Autonomous Database 13-5 13.6 Reference 14 14.1 **Terminology Information** 14-1 14.2 14-3 Installation Location 14.3 Log Files Location 14-3 14.4 Reset the Schema User Password 14-3 14.5 Use the OCI Console to Run Commands in the Data Transforms Instance 14-4 14.6 Increase the Memory of ODI Agent 14-5 14.7 Backup and Restore Data Transforms Repositories 14-5 Use SQL Developer to Export and Import Data Transforms Repositories 14-6 14.7.1 14.7.2 Switch to the Imported Repository 14-9 Update the Imported Repository using DataTransformsPostImport.sh 14.7.3 14-10



Preface

This guide helps you to get started and work with Oracle Data Transforms. This includes creating the stack, provisioning the repositories, accessing the instance, creating connections, loading data, and running data flows to transform your data.

This preface contains the following topics:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

This document is intended for administrators and users who want to deploy an Oracle Data Transforms instance and use it to load and transform data.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

For more information, see the following documents in *Oracle Data Integrator Library*:

- Release Notes for Oracle Data Integrator
- Understanding Oracle Data Integrator
- Developing Integration Projects with Oracle Data Integrator
- Installing and Configuring Oracle Data Integrator
- Upgrading Oracle Data Integrator
- Integrating Big Data with Oracle Data Integrator Guide



- Application Adapters Guide for Oracle Data Integrator
- Developing Knowledge Modules with Oracle Data Integrator
- Connectivity and Knowledge Modules Guide for Oracle Data Integrator Developer's Guide
- Oracle Data Integrator Tools Reference
- Data Services Java API Reference for Oracle Data Integrator
- Open Tools Java API Reference for Oracle Data Integrator
- Getting Started with SAP ABAP BW Adapter for Oracle Data Integrator
- Java API Reference for Oracle Data Integrator
- Getting Started with SAP ABAP ERP Adapter for Oracle Data Integrator
- Oracle Data Integrator 12c Online Help, which is available in ODI Studio through the JDeveloper Help Center when you press **F1** or from the main menu by selecting **Help**, and then **Search** or **Table of Contents**.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.



1

Overview

Learn about Oracle Data Transforms and its concepts.

Topics:

Introduction to Oracle Data Transforms

Oracle Data Transforms allows you to design graphical data transformations in the form of data loads, data flows, and workflows. Data Loads provide a convenient way of loading data into Oracle Autonomous Database.

Supported Browsers

Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR.

Supported Database Functions

Oracle Data Transforms supports various database functions that you can drag and drop on the Design Canvas to connect components within a data flow.

1.1 Introduction to Oracle Data Transforms

Oracle Data Transforms allows you to design graphical data transformations in the form of data loads, data flows, and workflows. Data Loads provide a convenient way of loading data into Oracle Autonomous Database.

Oracle Data Transforms provides an easy-to-use user interface to create your data loads, data flows, and workflows without requiring you to write any code. Data flows define how the data is moved and transformed between different systems, while the workflows define the sequence in which the data flows are executed. You can drag and drop the components, that is data entities and operators, such as Join, Split, and Aggregate, on the Design Canvas of the Data Flow editor and connect them to complete you data flow. You can then add your data flows to workflows.

After your data flows and workflows are ready, you can execute them immediately or schedule to execute them at a later time. Oracle Data Transforms run-time agent orchestrates the execution of jobs. On execution, Oracle Data Transforms generates the code for you.

Oracle Data Transforms integrates with Oracle Autonomous Database (Autonomous Data Warehouse and Autonomous Transaction Processing) for repository creation as well as for loading data. The Oracle Data Transforms Repository is automatically setup at provisioning. Objects developed or configured by the users are stored in this repository.

Oracle Data Transforms has many concepts similar to Oracle Data Integrator (ODI), however, the terminology used is different. If you are familiar with ODI, then reviewing the Terminology Information section may be helpful to you. See Terminology Information.

You can launch Data Transforms in any of the following ways:

 Oracle Cloud Marketplace: Create a Data Transforms instance from Oracle Cloud Marketplace. Data Transforms is available as a separate listing on Marketplace called Data Integrator: Web Edition. Autonomous Database Data Tools Page: Navigate to Autonomous Database
Data Tools Page, and click Data Transforms in the Database Actions page, or
click the Selector icon and select Data Transforms from the Data Tools menu in
the navigation pane.

If you have already registered a Data Transforms instance from OCI Marketplace with the Autonomous Database, the **Data Transforms** card on the Database Actions page will continue to take you to your Marketplace instance

Access to the standard set of Data Transforms features may depend on where you launch Data Transforms from. In this documentation, certain topics could include any of the following badges to indicate content that may or may not be available for use:

- APPLIES TO: Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.
- APPLIES TO: Data Transforms instance that is registered with Autonomous Database.

1.2 Supported Browsers

Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR.

Note that mobile browsers as well as private browsing mode is not supported for Firefox, Internet Explorer, or Edge.

1.3 Supported Database Functions

Oracle Data Transforms supports various database functions that you can drag and drop on the Design Canvas to connect components within a data flow.

The Database Functions toolbar in the Data Flow editor includes the following database functions that can be used in your data flows:

1. Data Transformation

It contains the following components:

- Aggregate
- Expression
- Filter
- Join
- Distinct
- Lookup
- Set
- Sort
- Subquery Filter
- Table Function

2. Data Preparation

It contains the following components:



- Data Cleanse
- Substitution
- Equi_Width Binning
- Quantile Binning
- Lead
- Lag
- Replace

3. Machine Learning

It contains the following components:

- Prediction
- Outlier Detection

4. Text

It contains the following components:

- REGEXP COUNT
- REGEXP INSTR
- REGEXP SUBSTR
- REGEXP REPLACE
- Edit Distance Similarity
- Contains

5. Oracle Spatial and Graph

It contains the following components:

- Buffer Dim
- Buffer Tol
- Distance Dim
- Distance Tol
- Nearest
- Simplify
- Point
- Geocode Tools:



The following Geocode Tools work only in non-Autonomous Database environment.

- Geocode As Geometry
- Geocode
- Geocode Addr
- Geocode All



- Geocode Addr All
- Reverse Geocode

Note:

The following Geocode Tool works only in an Autonomous Database environment.

- Geocode Cloud
- Spatial Join



Prerequisites

Make sure you have the following prerequisites before using Data Transforms on Oracle Cloud:

- Oracle Cloud Account
- Have access to assigned Oracle Cloud Tenant
- Compute node resources within Oracle Cloud Tenant

Go through the following prerequisites carefully before creating the Oracle Data Transforms instance:

- Refer to Create Dynamic Group and Policies for compartment, before provisioning Oracle Data Transforms and Autonomous Database Instances.
- Read the password policies mentioned in Manage Password Complexity on Autonomous Database. The password that you provide should conform to the password validation rules of the CLOUD_VERIFY_FUNCTION, which is the default password verification function for Autonomous Database.

Topics:

- Create an SSH/RSA Key
 - To work with the Oracle Cloud Infrastructure once the Oracle Data Transforms Compute Node is built, you have to provide a SSH Public Key to allow you to login to the node.
- Create Dynamic Group and Policies
 Dynamic Group and Policies are required when using the Oracle Data Transforms
 Repository on an Autonomous Database.

2.1 Create an SSH/RSA Key

To work with the Oracle Cloud Infrastructure once the Oracle Data Transforms Compute Node is built, you have to provide a SSH Public Key to allow you to login to the node.

In order to build your SSH keys, perform the following steps:

1. In a terminal window, generate the SSH key using the following command:

```
$ ssh-keygen Generating public/private rsa key pair.
```

2. Enter the path to store this file. By default, this gets saved in your home directory under a hidden folder called .ssh. Change this default location, if required.

Enter file in which to save the key (/Users/johndoe/.ssh/id rsa): <Return>

3. Enter a passphrase using your key.

```
Enter passphrase (empty for no passphrase): <passphrase>
```

4. Re-enter the passphrase to confirm it.

```
Enter same passphrase again: <passphrase>
```

5. Check the results.

The key fingerprint (a colon separated series of two digit hexadecimal values) is displayed. Check if the path to the key is correct. In the above example, the path is $/Users/johndoe/.ssh/id_rsa.pub$. You have now created a public and private key pair.

2.2 Create Dynamic Group and Policies

Dynamic Group and Policies are required when using the Oracle Data Transforms Repository on an Autonomous Database.

During instance provisioning, data servers for all accessible Autonomous Databases (ADB) are automatically created, as long as the dynamic group and policies are created before deploying Oracle Data Transforms on Marketplace and are set as mentioned below.

The provisioning wizard provides two options for you to select the autonomous database where the compute node will be built in. You need to create the dynamic group and policies depending on the option you want to use.

Search for the autonomous database using the OCID To use this option create a dynamic group to include the OCID of the compartment
that contains the database and the instance. For example,

```
ALL {instance.compartment.id = 'ocid1.compartment.oc1..aaaaaaaabgr34tpuanpvq6xfb667nsmy2jz45zj6dexojhxdsv4mjayem3cq'}
```

You have to setup the following policy at the compartment level:

```
Allow dynamic-group 
group-name> to inspect autonomous-database-
family in compartment <compartment-name>
Allow dynamic-group 
group-name> to read autonomous-database-family
in compartment <compartment-name>
Allow dynamic-group 
group-name> to inspect compartments in
compartment <compartment-name>
```

If the autonomous database and Data Transforms instance are in different subcompartments you need to create a dynamic group to include the OCID of the compartment that contains the instance. You then need to define a policy that grants the dynamic group access to the compartment that contains the database.

For example, let's assume

- RootCompartment
- SubcompartmentA contains the autonomous database



SubcompartmentB - contains the Data Transforms instance

In this case, if you want to search and provision using the OCID of SubcompartmentA you need to do the following:

 Create a dynamic group, say instance_grp, to include the OCID of SubcompartmentB. For example,

```
ALL {instance.compartment.id = 'OCID of SubcompartmentB'}
```

You then have to setup the following policy to allow <code>instance_grp</code> access to SubcompartmentA:

```
Allow dynamic-group instance_grp to inspect autonomous-database-family in compartment SubcompartmentA
Allow dynamic-group instance_grp to read autonomous-database-family in compartment SubcompartmentA
Allow dynamic-group instance_grp to manage instance-family in compartment SubcompartmentA
```

Browse to select the database -

To use this option create a dynamic group to include the compartment IDs of all the compartments within the tenancy (root compartment). For example,

```
ALL {instance.compartment.id =
'ocid1.compartment.oc1..aaaaaaaabgr34tpuanpvq6xfb667em3cnsmy2jz45zj6dexojh
xdsv4mjay4q'}
ALL {instance.compartment.id =
'ocid1.compartment.oc1..aaaaaaaa3elin667nsmpvq4fbytpua6x2jz4jayem35ojhxdsv
4mzj6dexcq'}
ALL {instance.compartment.id =
'ocid1.compartment.oc1..aaaaaaaaa7g4jr3pvq6xfb4jz45tpuan667nmsyz42j5zj6dexo
jhxdsv4mja'}
```

You have to setup the following policy at the tenancy (root compartment) level:

```
Allow dynamic-group 
group-name> to inspect autonomous-database-family in
compartment tenancy
Allow dynamic-group 
group-name> to read autonomous-database-family in
compartment tenancy
Allow dynamic-group 
group-name> to inspect compartments in compartment
tenancy
```

To configure email delivery service for specified groups on Oracle Cloud Marketplace:

An email approved sender must be in a group that has IAM policy permissions to send emails. An approved sender must be in a compartment with permissions to manage approved senders. You have to create a policy to manage approved senders in the entire tenant, if the approved senders exist in root compartment.

Add the following policy statement to enable odi group to manage approved senders:

Allow dynamic-group odi group to use approved-senders in compartment odi



3

Deploy

This section lists details about how to find Data Transforms on Oracle Cloud Marketplace, as well as steps to create the compute instance and provision the metadata repository.

Topics:

Select the Product

Oracle Data Transforms is listed on the Oracle Cloud Marketplace as Data Integrator: Web Edition.

- Create the Oracle Data Transforms Instance
 This topic contains step-by-step instructions for creating the Oracle Data Transforms instance.
- Add Ingress Rules to the Default Security List
 For establishing communication with the instance, you need to ensure that the underlying
 Security List that is associated with the VCN has all the IP protocols enabled.
- Provision the Oracle Data Transforms Metadata Repository
 The very first screen you will see after you access the Oracle Data Transforms URL is the
 Administration tab, which provides easy wizards that allow you to either create a new
 repository in an Autonomous Database to store your transformations, or connect to an
 existing Data Transforms repository.

3.1 Select the Product

Oracle Data Transforms is listed on the Oracle Cloud Marketplace as Data Integrator: Web Edition.

To search for the product Oracle Data Transforms on Oracle Cloud Marketplace, enter "Data Integrator: Web Edition" as the product name in the search text box and click **Go**. From the displayed listings click **Data Integrator: Web Edition**. You are navigated to the product listing page which has all the basic usage information along with the product documentation.



You may also see a listing called Data Integrator: Web Edition (Versions prior to May 2022), which includes older versions of Data Integrator: Web Edition. Do not use this listing. Oracle recommends that you use Data Integrator: Web Edition to get the latest version.

- 1. From your respective product listing page, select **Get App**.
- 2. Select OCI Region or Login using your Single Sign-On credentials.
 - OCI Region Select the desired region and click Launch Image.
- 3. Provide the OCI tenant details.
- Sign-in to the Identity provider.



- **5.** On the **Data Integrator: Web Edition** page, provide the following information:
 - Type It is Image by default.
 - **b.** Version Select a version from the drop-down. The latest version is selected by default.
 - c. Compartment Specifies the compartment where the compute node will be built. It is generally the location that you have access to build the compute node.
 - **d.** Terms of Use This check-box is selected by default. Oracle recommends to review the licenses before proceeding.
 - e. Launch Instance It launches the instance in the OCI environment. After selecting all the required information, click **Launch Instance**.



Provisioning your instance on Oracle Cloud Marketplace greatly depends on your product selection. You cannot use the same VM instance for multiple product offerings. The procedure listed below are product specific and may differ based on the product version.

3.2 Create the Oracle Data Transforms Instance

This topic contains step-by-step instructions for creating the Oracle Data Transforms instance.

Follow the below procedure to create your Oracle Data Transforms instance: Once you click **Launch Instance**, you are navigated to the **Create compute instance** page.

- **1.** Fill in the required information:
 - Name Name of the Instance. This field displays a default name based on the timestamp. You can edit this field, if required.
 - Create In Compartment This defaults to the compartment you have selected on the Oracle Data Integrator for Oracle page.
- 2. In the Placement settings, choose from any of the displayed availability domains. Do not change the values selected for **Capacity type** and **Fault Domain**.
- 3. In the Image and Shape settings, click **Change Shape** to select a template that you want to base the instance on.
- 4. In the Networking settings, fill in the required details to configure the networking options:
 - Primary network Do one of the following:
 - Choose Select exiting virtual cloud network, and then select the desired VCN for the newly created instance from the Virtual cloud network in compartment drop-down. A VCN is a software-defined network that you set up in the Oracle Cloud Infrastructure data centers in a particular region. Click Change Compartment if you want to change the VCN compartment where the compute node will be built.



 Choose Create new virtual cloud network, enter a name in New virtual cloud network name if you want to change the pre-populated value, and select the VCN compartment where the compute node will be built from the Create in compartment drop down.

You must add ingress rules to the default security list of the new created VCN to access the instance. See Add Ingress Rules to the Default Security List for more information.

- Choose Enter subnet OCID and enter the OCID of the compartment where the instance-related artifacts will be created in Subnet OCID.
- Subnet Do one of the following
 - Choose Select existing subnet and then select a subnet from the Subnet in compartment drop-down. If you want to change the compartment use Change Compartment to choose a different compartment.



You cannot use this option if you choose **Create new virtual cloud network** or **Enter subnet OCID** in the Primary network settings.

- Select Create new public subnet and enter the following details:
 - * New subnet name Enter a name if you want to change the pre-populated value.
 - * Create in compartment If you want to change the compartment select a different compartment from the drop-down.
 - * CIDR block Specify the Virtual Cloud Network (VCN) CIDR block for the VCN if you want to change the pre-populated value. You cannot change this value later.

Note:

You cannot see this option if you choose **Enter subnet OCID** in the Primary network settings.

- Public IP address Do one of the following:
 - Assign a public IPv4 address The newly created VM is assigned a public IP address.
 - Do not assign a public IPv4 address No public IP address is assigned to prevent public access to the compute node.
- 5. Add SSH keys so that you can connect to the instance using SSH. You can choose from the following options:
 - Generate a key pair for me Use this option to generate an SSH key pair to connect to the instance using an SSH connection.
 - Upload public key files (.pub) Use this option if you want to upload a .pub file that contains the SSH key pair.
 - Paste public keys Use this option if you want to manually specify the SSH key pair.
 - No SSH keys Use this option if you do not want to enable SSH access.



- 6. Leave the boot volume selection as is.
- 7. Click Create.
- **8.** When you get a message Successful from the OCI console, the instance is created. This takes approximately 5 minutes to complete.
- 9. After approximately 5 minutes, try to access the Data Transforms URL via a browser using the format http://<public-ip address>:9999/oracle-data-transforms. You can get IP address from the OCI Compute Instance properties page.

3.3 Add Ingress Rules to the Default Security List

For establishing communication with the instance, you need to ensure that the underlying Security List that is associated with the VCN has all the IP protocols enabled.

The following table lists the ingress rules you need to set to access your Data Transforms instance.

Table 3-1 Ingress Rules Table

Stateless	Source	IP Protocol	Source Port Range	Destinatio n Port Range	Type and Code	Allows
No	0.0.0.0/0	TCP	All	22	Nil	TCP traffic for ports: 22 SSH Remote Login Protocol
No	10.0.0.0/24	TCP	All	80	Nil	TCP traffic for ports : 80
No	0.0.0.0/0	All protocols	Nil	Nil	Nil	All traffic for all ports
No	0.0.0.0/0	ICMP	Nil	Nil	3, 4	ICMP traffic for: 3, 4 Destination Unreachabl e: Fragmentati on Needed and Don't Fragment was Set

3.4 Provision the Oracle Data Transforms Metadata Repository

The very first screen you will see after you access the Oracle Data Transforms URL is the Administration tab, which provides easy wizards that allow you to either create a



new repository in an Autonomous Database to store your transformations, or connect to an existing Data Transforms repository.

To provision the Oracle Data Transforms Repository on an Autonomous Database you need to create dynamic groups and policies depending on the option you want to use. Refer to Create Dynamic Group and Policies for details.

Follow the below procedure to provision the Oracle Data Transforms metadata repository:

- 1. On the Administration tab, select one of the following options:
 - Create a new Repository in an Autonomous Database
 - Connect to an existing Repository in an Autonomous Database

The repo mode created during configuration is marked as Oracle Data Transforms mode.

Note:

Oracle Data Integrator Repositories that you create during provisioning are mode specific. You cannot use a repository created in Oracle Data Transforms mode with a stack in ODI Studio mode and you cannot access a repository created in ODI Studio mode with a stack in Oracle Data Transforms mode.

- 2. In the Autonomous Database Details page, do one of the following:
 - Select Use OCID and then enter the OCID of the autonomous database that you
 want to use to search for the compartment where the compute node will be built in
 the database.

To use this option create a dynamic group to include the OCID of the compartment where the database instance is created, and then set up a policy to grant access to the dynamic group at the compartment level.

If the database and instance are in different subcompartments you need to create a dynamic group to include the OCID of the compartment that contains the instance. You then need to define a policy that grants the dynamic group access to the compartment that contains the database.

- Select Browse for Database and then do the following:
 - **a.** Autonomous Database Compartment Select the compartment where the compute node will be built in the database.
 - **b.** Autonomous Database Instance Select the created Autonomous Database instance from this drop down arrow.

To use this option create a dynamic group to include the compartment IDs of all the compartments within the tenancy (root compartment) and then set up a policy to grant access to the dynamic group at the tenancy (root compartment) level.

Refer to Create Dynamic Group and Policies for details about creating dynamic groups and policies for each option.

- 3. When you select the required Autonomous Database instance, Data Transforms generates a default password to download the wallet to create a connection to that instance. If you need to specify a custom password that conforms to specific validation rules, select the Use custom wallet password check box, and enter the password in the Wallet Password field.
- Click Fetch Service Names to retrieve the services associated with the database and select a name from the drop-down.



- 5. Provide the password credentials of the administrator user of the Autonomous Database in the Autonomous Database Admin Password textbox. The password that you provide should conform to the password validation rules of the CLOUD_VERIFY_FUNCTION, which is the default password verification function for Autonomous Database. For more information see, Manage Password Complexity on Autonomous Database.
- 6. In the Schema Details page, enter the following:
 - **Schema Prefix** Specify the prefix used for the Oracle Data Transforms Schemas to identify them.



A valid prefix should contain only alpha-numeric characters. It should not start with a number and should not contain any special characters. A maximum of 12 characters are allowed.

For existing repository, the Schema Prefix field is a drop down that lists the schema prefixes available based on the Autonomous Database Instance you selected in the Autonomous Database Details page.

 Schema Password - Specify the password used for accessing the Oracle Data Transforms schema.

Note:

The password that you provide should conform to the password validation rules of the CLOUD_VERIFY_FUNCTION, which is the default password verification function for Autonomous Database. For more information see, Manage Password Complexity on Autonomous Database.

For existing repository, provide the password configured during repository creation. Note that the schema password expires in 180 days from the date when the repository was created. In such cases, provisioning fails when you try to connect to the existing repository. See Reset the Schema User Password for instructions on how to reset the schema password.

7. In the Repository Details page, provide the ODI Supervisor password.

Note:

The Supervisor password that you provide should:

- contain alpha-numeric characters,
- be 6-12 characters in length,
- not start with a number,
- not contain any special characters except \$, # and _.



- 8. Click **Save**. The final screen of the wizard displays all the processes that run to complete the configuration. This takes about 15 minutes. The screen auto-refreshes to display the status of each process. You can also click the Refresh icon to view the status of the configuration.
- 9. Click **Download Log File** to download the configuration logs. The log files are downloaded in a zip file named provlogs.zip.
- **10.** When the configuration completes successfully, the **Login to Data Transforms** button is activated, which allows you to navigate to the Oracle Data Transforms login page.

If the configuration fails at any step you will not be able to see the login page. To see the status of the configuration, follow the instructions to connect the Data Transforms Instance as described in Connect to Oracle Data Transforms Instance. Check the transformsStartup.log file, fix the issues, and repeat these steps to provision the repository.

If you want to check the transformsStartup.log file without connecting to the instance, run the displayStartupLog command from the OCI Console using the compute instance run command feature. See Use the OCI Console to Run Commands in the Data Transforms Instance for detailed instructions.



4

Access

This section includes information about the many ways you can use to connect to the Data Transforms compute instance.

Topics:

- Connect to Oracle Data Transforms Instance
 This section includes information about connecting to the Data Transforms compute instance by using a Secure Shell (SSH) connection.
- Access the Oracle Data Transforms Compute Instance through VNC
 To access the Oracle Data Transforms instance through VNC, you need to install a VNC viewer on your local computer. You also need to provide the VNC password that you had provided during the stack creation.
- Access Oracle Data Transforms
 This topic provides information on how you can access the Oracle Data Transforms from a web browser.

4.1 Connect to Oracle Data Transforms Instance

This section includes information about connecting to the Data Transforms compute instance by using a Secure Shell (SSH) connection.



This section is applicable only for Advanced administrative users who would typically connect to the Data Transforms instance to examine the logs for errors. If you are a general Data Transforms developer, you can skip this section.

Topics:

- Connect to Oracle Data Transforms Instance from Linux
- Connect to Oracle Data Transforms Instance from Windows

4.1.1 Connect to Oracle Data Transforms Instance from Linux

You can connect to an Oracle Data Transforms compute instance by using a Secure Shell (SSH) connection. Most Linux distributions include an SSH client by default.

To connect to your Oracle Data Transforms compute instance from Linux,

1. To enable SSH access to Oracle user, log in to your instance using SSH as opc user.

ssh opc@<IP Address>

Execute the following commands to copy the authorized keys to Oracle user:

```
sudo cp /home/opc/.ssh/authorized_keys /home/oracle/.ssh
sudo chown oracle:oracle /home/oracle/.ssh/authorized keys
```

Use the following command to set the file permissions so that only you can read the file:

```
$ chmod 400 <private key>
```

where <pri>private_key> is the full path and name of the file that contains the private key associated with the instance you want to access.

3. Use the following SSH command to access the instance.

```
$ ssh -i <private key> <username>@<public-ip-address>
```

where

- <private_key> is the full path and name of the file that contains the private key associated with the instance you want to access.
- <username> is the default name for the instance. The default user name is oracle.
- <public-ip-address> is your instance IP address that you retrieved from the Console.

4.1.2 Connect to Oracle Data Transforms Instance from Windows

You can connect to an Oracle Data Transforms compute instance by using a Secure Shell (SSH) connection. For Windows, you can download a free SSH client called PuTTY from http://www.putty.org.

To connect to your Oracle Data Transforms compute instance from Windows,

- 1. Open putty.exe.
- 2. In the Category pane, select Window, and then select Translation.
- 3. In the Remote character set drop-down list, select UTF-8. The default locale setting on Linux-based instances is UTF-8, and this configures PuTTY to use the same locale.
- 4. In the Category pane, select Session and enter the following:
 - a. Host Name (or IP address):<username>@<public-ip-address>, where <username> is the default name for the instance. For Oracle Linux and CentOS images, the default user name is oracle. For the Ubuntu image, the default name is ubuntu and <public-ip-address> is your instance public IP address that you retrieved from the console.
 - b. Port: 22
 - c. Connection type: SSH
- 5. In the Category pane, expand Connection, expand SSH, and then click Auth.
- 6. Click Browse, and then select your private key.



7. Click Open to start the session.

If this is your first time connecting to the compute instance, you might see a message that the server's host key is not cached in the registry. Click **Yes** to continue the connection.

4.2 Access the Oracle Data Transforms Compute Instance through VNC

To access the Oracle Data Transforms instance through VNC, you need to install a VNC viewer on your local computer. You also need to provide the VNC password that you had provided during the stack creation.

To access Oracle Data Transforms utilities and logs through VNC, do the following:

 Log in to the provisioned Oracle Data Transforms instance on Oracle Cloud Marketplace using SSH as opc user:

```
ssh -i <private key> opc@<IP Address>
```

2. Execute the following firewall commands to open the VNC ports:

```
sudo firewall-cmd --permanent --new-service=odissh
sudo firewall-cmd --permanent --service=odissh --set-description="ODI
SSHserver"
sudo firewall-cmd --permanent --service=odissh --add-port=5901-5905/tcp
sudo firewall-cmd --permanent --add-service=odissh
sudo firewall-cmd -reload
```

3. To create a VNC session, log in to the instance as Oracle user:

```
ssh -i <private_key> oracle@<IP Address>
```

- 4. Start the VNC server by typing vncserver.
- **5.** Enter the VNC password that you had provided during the stack creation.
- 6. Install a VNC viewer on your local computer.
- Use SSH to connect to the compute instance running Oracle Data Transforms, as described in Connect to Oracle Data Transforms Instance.
- 8. On your local computer, connect to your instance and create a ssh tunnel for port 5901 (for display number 1):

```
$ ssh -L 5901:localhost:5901 -i id rsa oracle@<IP Address>
```

9. On your local computer, for the VNC to work, add an Ingress rule as follows:

```
No 0.0.0.0/15 TCP All 5901 TCP traffic for ports: 5901
```

- **10.** On your local computer, start a VNC viewer and establish a VNC connection to localhost:1.
- **11.** Enter the VNC password that you had provided during the stack creation.



12. For connecting multiple users, after the vncpasswd utility exits, start the VNC server by typing vncserver. This will start a VNC server with display number 1 for the oracle user, and the VNC server starts automatically if your instance is rebooted. For example vncserver@:2 or vncserver@:3.

Note:

Do not reboot the instance by choosing **Force reboot the instance by immediately powering off**. The VNC session will not start automatically.

To start developing your data loads, data flows or workflows,

- 1. In a web browser that has network connectivity to your compute instance, enter the url http://<ip address>:9999/oracle-data-transforms.
- 2. Sign in to Data Transforms using the user and password supplied during instance deployment. Default user is SUPERVISOR.
- 3. Click **Connections** from the left pane of the **Home** page.
- **4.** Click **Test connection**, to check if the created Autonomous Database connection is working.

4.3 Access Oracle Data Transforms

This topic provides information on how you can access the Oracle Data Transforms from a web browser.

You can access Oracle Data Transforms interface on Oracle Cloud Marketplace by:

- Navigating to Autonomous Database Data Tools Page
 For more details, refer to Access Oracle Data Transforms From Database Actions page.
- Click the link in OCI Stack Details page available in the Application Information tab.
- In a web browser that has network access to the Oracle Data Transforms
 Instance, type in <ip address>:9999/oracle-data-transforms/ in the toolbar and click Enter.

You are navigated to the login page.

In the **Sign Into Oracle Data Integrator** page, enter the following login credentials and click **Connect**:

- Username SUPERVISOR
- Password Use the password that you provided during the ODI stack provisioning.

The Oracle Data Transforms interface appears allowing you to perform the data transformations. The very first screen you will see after you log in is the Home tab, which provides easy wizards to complete the following two actions:

Load Data – Allows you to load multiple tables from the source connection schema
to the target connection schema. You can also do this by clicking the required
project tile in the Projects tab, clicking Data Loads, and then clicking Create Data
Load.



Transform Data - Allows you to create data flows to map data and run transformations.
 You can also do this by selecting an existing project and on the Project Details page, clicking Create Data Flow.



5

Connections

Learn about the connection types that Oracle Data Transforms supports and how to work with connections. This section also includes topics that cover instructions that are specific to certain connections.

Topics:

Work with Connections

Connections help you to connect Data Transforms to various technologies reachable from your OCI network.

Supported Connection Types

This topic lists the connection types that are supported for connecting to Data Transforms.

- Create Custom Connectors
- Create a Data Transforms Connection for Remote Data Load
 You can connect to an existing Data Transforms instance and run a data load remotely.
- Create an Oracle Business Intelligence Cloud Connector Connection
 Oracle Business Intelligence Cloud Connector (BICC) allows you to extract business data from a data source and load it into configured external storage.
- Create a Delta Share Connection

Databricks Delta Share is an open protocol for secure data sharing. Oracle Data Transforms integrates with Delta Share to load data to Oracle Autonomous Database. You can use the Delta Share connection to load data from Databricks or Oracle Data Share.

Create an Oracle Financials Cloud Connection

You can fetch real time transactional data from Oracle Financials Cloud REST endpoints, import the data entities into Data Transforms, and use them as a source in a data flow.

Create and Use an Oracle NetSuite Connection

Data Transforms uses the Oracle NetSuite JDBC Driver to connect to the Oracle NetSuite application. For Oracle NetSuite connections, Data Transforms allows you to load pre-built dataflows and workflows that you can run to transfer data from NetSuite to your target schema.

- Create an Oracle Object Storage Connection
 You can use Data Transforms for uploading data from Oracle Object Storage to
 Autonomous Database.
- Create a REST Server Connection

5.1 Work with Connections

Connections help you to connect Data Transforms to various technologies reachable from your OCI network.

This section describes the generic steps to create a connection. The displayed connection detail options may vary depending on the selected connection type.

Apart from the connection types listed in Supported Connection Types you can create custom connectors, which you can use to connect Data Transforms to any JDBC supported data sources. See Create Custom Connectors.

To create a new connection:

1. From the left pane of the Home page, click the **Connections** tab.

Connections page appears.

2. Click Create Connection.

Create Connection page slides in.

- 3. Do one of the following:
 - In the **Select Type** field, enter the name or part of the name of the connection type.
 - Select the type of connection that you wish to create.
 - Databases Allows you to configure any connection type for supported database types.
 - Applications Allows you to configure any connection type for supported applications.
 - Services Allows you to configure any connection type for supported services.
- 4. After selecting the required connection type, click **Next**.
- The Connection Name field is pre-populated with a default name. You can edit this value.
- **6.** For **Connection Details**, provide the connection details for the selected type such as:
 - Connection -
 - JDBC URL The URL to connect to the data server. For example: jdbc:weblogic:sqlserver://hostname:port[;property=value[;...]]

For connectors that use an autoREST driver that provides the model files along with the driver, specify the servername and other properties required to connect to that datasource. For example:

```
jdbc:weblogic:autorest://servername;[property=value[;...]]
```

- User The user name, if required, for connecting to the server.
- Password The password for connecting to server.
- Advanced Options
 - Array Fetch Size When reading large volumes of data from a data server,
 Oracle Data Transforms fetches successive batches of records. This value is the number of rows (records read) requested by Oracle Data
 Transforms on each communication with the data server.
 - Batch Update Size When writing large volumes of data into a data server, Oracle Data Transforms pushes successive batches of records.
 This value is the number of rows (records written) in a single Oracle Data Transforms INSERT command.



 Degree of Parallelism for Target - This value indicates the number of threads allowed for a loading task. The default value is 1. The maximum number of threads allowed is 99.



Connection details are specific and the above options vary based on the selected connection type.

After providing all the required connection details, click Test Connection to test the connection.

If the test connection fails, check whether the Autonomous Database from where you are accessing Data Transforms is configured to use a private endpoint. See Enable Access to Private Data Sources from Autonomous Database for more information.

8. Click Create.

The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (*) next to the selected connection to perform the following operations:

- Select Edit to edit the provided connection details.
- Select Test Connection to test the created connection.
- Select Delete Schema to delete schemas.
- Select Delete Connection to delete the created connection.

You can also search for the required connection to know its details based on the following filters:

- Name of the connection.
- Technology associated with the created connection.

5.2 Supported Connection Types

This topic lists the connection types that are supported for connecting to Data Transforms.



APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

- For the connectors that require driver installation, you need to copy the jar files to the /u01/oracle/transforms_home/userlibs directory before you add the connection.
- Apart from the connection types listed here you can create custom connectors, which you can use to connect Data Transforms to any JDBC supported data sources. See Create Custom Connectors.



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Aha!	Application	Yes	Yes	Yes	No	
Ahrefs	Application	Yes	Yes	Yes	No	
Amazon Aurora	Database	Yes	Yes	Yes	Yes	
Amazon EMR Hive	Database	Yes	Yes	Yes	No	
Amazon Redshift	Database	Yes	Yes	Yes	Yes	
Apache Hive	Database	Yes	Yes	Yes	No	
Apache Spark SQL	Database	Yes	Yes	Yes	No	
AWS S3	Database	Yes	Yes	Yes	No	
Azure Billing	Application	Yes	Yes	Yes	No	
Azure Compute	Database	Yes	Yes	Yes	No	
Azure Data Lake Storage	Database	Yes	Yes	Yes	No	
Azure Reserved VM Instances	Database	Yes	Yes	Yes	No	
Azure Resource Health	Application	Yes	Yes	Yes	No	
Azure SQL Database	Database	Yes	Yes	Yes	Yes	
Azure Synapse Analytics	Database	Yes	Yes	Yes	Yes	
BigCommer ce	Application	Yes	Yes	Yes	No	Requires driver installation
Cassandra	Database	Yes	Yes	Yes	Yes	
Cloudera CDH Hive	Database	Yes	Yes	Yes	No	
Cloudera Impala	Database	Yes	Yes	Yes	No	
Confluence Cloud	Database	Yes	Yes	Yes	No	



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Data Transforms	Service	Yes	Yes	Yes	No	For instructions on connecting to an existing Data Transforms instance, see Create a Data Transforms Connection for Remote Data Load.
DataStax	Application	Yes	Yes	Yes	Yes	
Delta Share	Application	Yes	Yes	Yes	No	For instructions on creating a connection using Delta Share, see Create a Delta Share Connection
DocuSign	Database	Yes	Yes	Yes	No	
еВау	Application	Yes	Yes	Yes	No	Requires driver installation
EnterpriseD B	Database	Yes	Yes	Yes	Yes	
FinancialFor ce	Application	Yes	Yes	Yes	Yes	
FourSquare	Application	Yes	Yes	Yes	No	
Generic Rest	Application	Yes	Yes	Yes	No	For information about connecting to any REST service endpoint to create a connection, see Create a REST Server Connection.



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Generic Rest Config	Application	Yes	No	No	No	For information about connecting to any REST service endpoint to create a connection, see Create a REST Server Connection.
GitHub	Application	Yes	Yes	Yes	No	
Google Ads	Application	Yes	No	No	Depends on the driver	Requires driver installation
Google AdSense	Application	Yes	Yes	Yes	No	
Google Analytics	Application	Yes	Yes	Yes	No	
Google BigQuery	Database	Yes	Yes	Yes	No	
Google Calendar	Application	Yes	Yes	Yes	No	
Google Campaign Manager	Application	Yes	Yes	Yes	No	
Google Contacts	Application	Yes	Yes	Yes	No	
Google Drive	Database	Yes	Yes	Yes	No	
Google Search Ads 360	Application	Yes	Yes	Yes	No	
Greenplum	Database	Yes	Yes	Yes	No	
Hortonwork s Hive	Database	Yes	Yes	Yes	No	
HubSpot	Application	Yes	Yes	Yes	No	
Hypersonic SQL	Database	Yes	Yes	Yes	Yes	
IBM BigInsights	Database	Yes	Yes	Yes	No	
IBM DB2 Hosted	Database	Yes	Yes	Yes	Yes	



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
IBM DB2 UDB	Database	Yes	Yes	Yes	Yes	
IBM DB2 Warehouse	Database	Yes	Yes	Yes	Yes	
IBM DB2/400	Database	Yes	Yes	Yes	Yes	
Informix	Database	Yes	Yes	Yes	No	
Jira	Application	Yes	Yes	Yes	No	
Klaviyo	Application	Yes	Yes	Yes	No	
Magento	Application	Yes	No	No	Depends on the driver	Requires driver installation
Mailchimp	Application	Yes	Yes	Yes	No	
MapR Hive	Database	Yes	Yes	Yes	No	
Marketo	Application	Yes	Yes	Yes	No	
Microsoft Dynamics 365	Application	Yes	Yes	Yes	Yes	
Microsoft SharePoint	Application	Yes	Yes	Yes	Yes	
Microsoft SQL Server	Database	Yes	Yes	Yes	Yes	
Mongo DB	Database	Yes	Yes	Yes	Yes	
MySQL	Database	Yes	Yes	Yes	Yes	Make sure that the system variable property sql_require_primary_key is set to OFF. Otherwise, an ADW to MySQL mapping could fail with a "Table does not exist"



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
MySQL Heatwave	Database	Yes	Yes	Yes	Yes	Make sure that the system variable property sql_require_primary_key is set to OFF. Otherwise, an ADW to MySQL Heatwave mapping could fail with a "Table does not exist" error.
Netezza	Database	Yes	No	No	Depends on the driver	Oracle Data Transforms uses the Netezza JDBC to connect to a NCR Netezza database. This driver must be installed in your Data Transforms userlibs directory. See Download the Netezza JDBC driver for more information.



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Oracle	Database	Yes	Yes	Yes	Yes	For Data Integrator Web Edition, write operation is supported only on Oracle cloud database targets. For details refer to the Oracle terms of use before deploying the image from OCI marketplace
Oracle Analytics Cloud	Application	Yes	Yes	Yes	No	
Oracle Business Intelligence Cloud (BICC) Connector	Application	Yes	Yes	Yes	No	For information about creating a connection using Oracle Business Intelligence Cloud (BICC) Connector, see Create an Oracle Business Intelligence Cloud Connector Connection.
Oracle EBS	Application	Yes	Yes	Yes	Yes	



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Oracle Financials Cloud	Application	Yes	Yes	Yes	No	For information about creating a connection using Oracle Financials Cloud, see Create an Oracle Financials Cloud Connection.
Oracle Fusion ERP	Application	Yes	Yes	Yes	No	
Oracle Fusion Sales	Application	Yes	Yes	Yes	No	
Oracle Fusion Service	Application	Yes	Yes	Yes	No	
Oracle GoldenGate – OCI	Service	Yes	Yes	Yes	Yes	
Oracle Marketing Cloud	Application	Yes	Yes	Yes	Yes	
Oracle NetSuite	Application	Yes	Yes	Yes	No	For information about creating a connection using Oracle Netsuite, see Create and Use an Oracle NetSuite Connection.



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Oracle Object Storage	Database	Yes	Yes	Yes	Yes	For information about creating a connection using Oracle Object Storage, see Create an Oracle Object Storage Connection.
Oracle People Soft	Application	Yes	Yes	Yes	No	
Oracle Sales Cloud	Application	Yes	Yes	Yes	No	
Oracle Service Cloud	Application	Yes	Yes	Yes	No	
Oracle SIEBEL	Application	Yes	Yes	Yes	No	
PayPal	Application	Yes	Yes	Yes	No	
Pivotal HD	Database	Yes	Yes	Yes	No	
Pivotal HDB	Database	Yes	Yes	Yes	No	
PostgreSQL	Database	Yes	Yes	Yes	Yes	
Qmetry	Application	Yes	Yes	Yes	No	
QuickBooks Online	Application	Yes	Yes	Yes	No	
QuickBooks Payments	Application	Yes	Yes	Yes	No	
Quora Ads	Application	Yes	Yes	Yes	No	
Sage	Application	Yes	Yes	Yes	No	
Salesforce Chatter	Application	Yes	Yes	Yes	No	
Salesforce.c om	Application	Yes	Yes	Yes	Yes	
SAP BW/ 4HANA	Database	Yes	Yes	Yes	No	
SAP HANA	Application	Yes	Yes	Yes	No	
SAP NetWeaver	Database	Yes	Yes	Yes	No	
SAP S/ 4HANA Cloud	Application	Yes	Yes	Yes	No	



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Semrush	Application	Yes	Yes	Yes	No	
ServiceNow	Service	Yes	Yes	Yes	No	
Shopify	Application	Yes	Yes	Yes	No	Requires driver installation
Snowflake	Database	Yes	Yes	Yes	Yes	
Square	Application	Yes	Yes	Yes	No	
Stripe	Application	Yes	Yes	Yes	No	
Sybase As Anywhere	Database	Yes	Yes	Yes	Yes	
Sybase as Enterprise	Database	Yes	Yes	Yes	Yes	
Sybase AS IQ	Database	Yes	Yes	Yes	Yes	
TeamCity	Application	Yes	Yes	Yes	No	
Teradata	Database	Yes	No	No	Depends on the driver	Data Transforms uses the Teradata JDBC Driver to connect to a Teradata Database. To use Teradata as a data source the Teradata Gateway for JDBC must be running, and this driver must be installed in your Data Transforms userlibs directory. You will find the driver here: https:// downloads.t eradata.co m/ download/ connectivity/ jdbc-driver.



Name	Туре	Supported in Data Integrator: Web Edition	Supported in Data Transform s built into Autonomo us Database	Supported in Data Transform s built into OCI GoldenGat e	Supports Write Operation	Notes
Tumblr	Application	Yes	Yes	Yes	No	
Twitter	Application	Yes	Yes	Yes	No	
Veeva CRM	Application	Yes	Yes	Yes	Yes	
Volusion	Application	Yes	Yes	Yes	No	
Wistia	Application	Yes	Yes	Yes	No	
WooComme rce	Application	Yes	No	No	Depends on the driver	Requires driver installation
WordPress	Application	Yes	Yes	Yes	No	
Workday	Application	Yes	No	No	Depends on the driver	Requires driver installation
Xero	Application	Yes	Yes	Yes	No	Requires driver installation
Yelp	Application	Yes	Yes	Yes	No	
Zendesk	Application	Yes	Yes	Yes	No	Requires driver installation
Zoho CRM	Application	Yes	Yes	Yes	No	
Zoom	Application	Yes	Yes	Yes	No	

5.3 Create Custom Connectors

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

The Custom Connections page of the Administration tab of Oracle Data Transforms helps you to create custom connectors that point to any JDBC supported data sources.

The custom connectors will be listed in the Create Connection page where you can use them to connect data sources to Data Transforms. See Work with Connections for more information.

To create a new connector:

- In the left pane, click Administration.
 A warning message appears.
- 2. Click Continue.
- **3.** In the left pane, click **Custom Connections**. Custom Connections screen appears.
- 4. Click Create Connection Type.
 The Create Connection Type page appears.



- **5.** From the **Category** drop-down select the type of connection that you wish to create whether database, application, or service.
- 6. Enter a name for the connection.
- Enter the name of the JDBC Driver.
- 8. Click OK.

The newly created custom connection appears in the list and are available in the Create Connection page.

Click the ${\bf Actions}$ icon (*) next to the selected connection to perform the following operations:

- Select Edit, to edit the provided connection details.
- Select **Delete**, to delete the created connection.



You cannot delete custom connectors that have existing connections.

5.4 Create a Data Transforms Connection for Remote Data Load

You can connect to an existing Data Transforms instance and run a data load remotely.

To create this connection, you need to specify the URL of the Data Transforms instance along with the name of the ODI rest API from where you want to run the data load.

To define a Data Transforms connection:

- **1.** From the left pane of the Home page, click the **Connections** tab.
 - **Connections** page appears.
- 2. Click Create Connection.

Create Connection page slides in.

- 3. For Select Type,
 - In the Name field, enter the name of the newly created connection
 - Select Services as the type of connection that you wish to create.
- 4. In the **Endpoint URL** textbox, enter the URL of the ODI rest API from where you want to run the data load. Enter the URL in the format http://<host-ip-address>:<port>/odi-rest.
- 5. In the **User** text box enter SUPERVISOR as the user name.
- 6. In the **Password** text box enter the ODI Supervisor password.
- 7. After providing all the required connection details, click **Test Connection** to test the established connection.



Click Create.

The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (*) next to the selected connection to perform the following operations:

- Select Edit to edit the provided connection details.
- Select Test Connection to test the created connection.
- Select Delete Schema to delete schemas.
- Select **Delete Connection** to delete the created connection.

You can also search for the required connection to know its details based on the following filters:

- Name of the Connection.
- Technology associated with the created Connection.

5.5 Create an Oracle Business Intelligence Cloud Connector Connection

Oracle Business Intelligence Cloud Connector (BICC) allows you to extract business data from a data source and load it into configured external storage.

To create an Oracle BICC connection you need to first configure external storage using the OCI Object Storage Connection tab in the BICC Console. You need to specify these connection details when you define the connection in Oracle Data Transforms.

You can use the BICC connection to choose the offerings whose data stores you want to extract. Data Transforms uses an Oracle Object Storage Data Server used by Oracle BICC to stage the extracted files, which you can then use as a source for mapping. Note that you cannot use an Oracle BICC connection as a target for mapping.

To define an Oracle BICC connection,

- 1. From the left pane of the Home page, click the **Connections** tab.
 - Connections page appears.
- 2. Click Create Connection.

Create Connection page slides in.

- 3. Do one of the following:
 - In the Select Type field, enter the name or part of the name of the connection type.
 - Select the Applications tab.
- 4. Select **Oracle BI Cloud Connector** as the connection type.
- 5. Click Next.
- The Connection Name field is pre-populated with a default name. You can edit this value.
- Enter the URL in the BI Cloud Connector Service URL textbox.



- **8.** In the **Connection** section, enter the following details:
 - In the **User** text box enter your Oracle Cloud Infrastructure username.
 - In the Password text box enter the auth token.
- 9. In the **Storage** section, enter the following details:
 - In the **External Storage BICC Name** text box enter the name of the external storage as it appears in the Oracle BI Cloud Connector Console.
 - In the External Storage Bucket text box specify the bucket into which extracts are uploaded. Bucket names are obtained in the OCI Console.
 - In the External Storage Name Space text box specify the namespace.
 Namespace is obtained in the OCI Console.
 - In the External Storage Region text box enter the region.
 - In the External Storage User text box enter the name of the user logging into the Oracle Storage Cloud Service. The value specified in this field must be the same as the User Name configured in the Oracle BI Cloud Connector Console.
 - In the External Storage Token text box enter the password of the logged in user. The value specified in this field must be the same as the Password configured in the Oracle BI Cloud Connector Console.
- 10. Click **Test Connection** to test the established connection.
- 11. Click Create.

The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the $\bf Actions$ icon ($\ ^*$) next to the selected connection to perform the following operations:

- Select Edit, to edit the provided connection details.
- Select Test Connection, to test the created connection.
- Select Delete Schema, to delete schemas.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- Name of the Connection.
- Technology associated with the created Connection.

5.6 Create a Delta Share Connection

Databricks Delta Share is an open protocol for secure data sharing. Oracle Data Transforms integrates with Delta Share to load data to Oracle Autonomous Database. You can use the Delta Share connection to load data from Databricks or Oracle Data Share.

To use Databricks as a source, you need to specify the URL of the Delta Sharing server along with the bearer token that lets you access the Delta Lake share server. To



use Oracle Data Share as a source, you need to specify the URL for the token end point along with a client ID and the secret key.

This topic has the following sections:

- Creating the Delta Share Connection
- Creating and Running a Delta Share Data Load

Creating the Delta Share Connection

To define a Delta Share connection:

- From the left pane of the Home page, click the Connections tab.
 Connections page appears.
- 2. Click Create Connection.
 Create Connection page slides in.
- 3. Do one of the following:
 - In the Select Type field, enter the name or part of the name of the connection type.
 - Select the **Databases** tab.
- 4. Select **Delta Share** as the connection type.
- Click Next.
- The Connection Name field is pre-populated with a default name. You can edit this value.
- 7. In the **Share Endpoint URL** textbox, enter the URL of the Delta Sharing server. Enter the value in the <host>:<port>/<shareEndpoint>/ format.
- 8. In the Connection section, do one of the following:
 - Select Oracle Data Share and provide the Token Endpoint URL, Client ID, and Client Secret for accessing the share.
 - You can get this information from the Delta Share Profile JSON document that you will need to download from supplied to you by the Share Provider. (This is also where they get the Share Endpoint URL from)
 - You can get this information from the Delta Share Profile JSON document that you can download from the activation link that is provided by the Data Share provider to access their share.
 - Select Databricks and in the Bearer Token text box enter the token for connecting to the Delta Sharing server.
- 9. If you need to use a proxy to access the Delta Share Server or Delta Share Storage configure the following settings:
 - In the Proxy Host textbox, enter the host name of the proxy server to be used for the connection.
 - In the Proxy Port textbox, enter the port number of the proxy server.
 - Select the following checkboxes depending on where the proxy is required:
 - Use Proxy to access Delta Share Server
 - Use Proxy to access Delta Share Storage
- **10.** Click **Test Connection**, to test the established connection.
- 11. After providing all the required connection details, click **Create**.



The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (*) next to the selected connection to perform the following operations:

- Select Edit, to edit the provided connection details.
- Select Test Connection, to test the created connection.
- Select Delete Schema, to delete schemas.
- Select Delete Connection, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- Name of the Connection.
- Technology associated with the created Connection.

Creating and Running a Delta Share Data Load

To load data from Delta Share into Oracle Autonomous Database, the Oracle connection user must be an Admin user. Admin privileges are required so that the Oracle user can create and insert data into tables in another schema.

When you run the data load, Data Transforms loads the data onto a corresponding table in the target schema. The data load runs incrementally. The very first time you run a data load, Data Transforms copies all the data into new tables. For every subsequent data load run, it only uploads the changes. Any additions or deletions in the records will reflect in the target tables. Note that if there is any metadata change in the table, for example a column is added, Data Transforms creates a new table to load the data on to the target server. You could create a workflow, add the data load as a step, create a schedule to run the workflows at a predefined time interval. See

To create and run a Delta Share data load:

- 1. Do one of the following:
 - On the Home page, click Load Data. The Create Data Load wizard appears.
 In the Create Data Load tab, enter a name if you want to replace the default value, add a description, and select a project from the drop-down.
 - On the Home page, click **Projects**, and then the required project tile. In the left pane, click **Data Loads**, and then click **Create Data Load**. The Create Data Load wizard appears.
 - In the Create Data Load tab, enter a name if you want to replace the default value and add a description.
- Click Next.
- 3. In the Source Connection tab,
 - a. From the **Connection Type** drop-down, select **Delta Share**.
 - **b.** from the **Connection** drop-down, select the required connection from which you wish to add the data entities.
 - **c.** Select the share that you want to load tables from the **Share** drop-down. The drop-down lists all the shares for the selected connection.



- d. Click Next.
- 4. In the Target Connection tab,
 - a. From the **Connection Type** drop-down, select **Oracle** as the connection type.



This drop-down lists only JDBC type connections.

- **b.** From the **Connection** drop-down, select the required connection from to you wish to load the data entities.
- c. Enter a unique name in the **Schema** textbox.
- d. Click Save.

The Data Load Detail page appears listing all the tables in the selected share with their schema names.



For Delta Share data loads the Data Load Detail page only includes the option. You cannot apply different actions - incremental merge, incremental append, recreate, truncate, append - on the data entities before loading it to the target schema. This is to make sure that the data is consistent between the Delta Sharing server and the target schema.

5. Click to run the data load.
A confirmation prompt appears when the data load starts successfully.

To check the status of the data load, see the Status panel on the right below the Target Schema details. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows links to the jobs that execute to run this data load. Click the link to monitor the progress on the Job Details page. For more information about jobs, see Create and Manage Jobs.

All the loaded data entities along with their details are listed in the Data Entities page. To view

the statistics of the data entities, click the **Actions** icon (*) next to the data entity, click **Preview**, and then select the **Statistics** tab. See View Statistics of Data Entities for information.

5.7 Create an Oracle Financials Cloud Connection

You can fetch real time transactional data from Oracle Financials Cloud REST endpoints, import the data entities into Data Transforms, and use them as a source in a data flow.

To create an Oracle Financials Cloud connection you need to choose a temporary schema where Data Transforms can create data entities after the reverse-engineering operation.

To define an Oracle Financials Cloud connection,

1. From the left pane of the Home page, click the **Connections** tab.



Connections page appears.

2. Click Create Connection.

Create Connection page slides in.

- 3. Do one of the following:
 - In the **Select Type** field, enter the name or part of the name of the connection type.
 - Select the Applications tab.
- 4. Select **Oracle Financials Cloud** as the connection type.
- 5. Click Next.
- **6.** The **Connection Name** field is pre-populated with a default name. You can edit this value.
- In the REST Service URL textbox, enter the URL of the endpoint that services the REST resources.
- 8. In the **Proxy Host** textbox, enter the host name of the proxy server to be used for the connection.
- 9. In the **Proxy Port** textbox, enter the port number of the proxy server.
- 10. In the **User** text box enter the user name for connecting to the REST endpoint.
- 11. In the **Password** text box enter the password for connecting to the REST endpoint.
- **12.** Choose a connection from the **Staging Connection** drop-down list. The list displays only existing Autonomous Database connections. To use a different connection, create the connection before you reach this page.
- 13. After providing all the required connection details, click **Create**.

The new connection is created.

14. Click **Test Connection**, to test the established connection.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (*) next to the selected connection to perform the following operations:

- Select Edit, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Select **Delete Schema**, to delete schemas.
- Select Delete Connection, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- Name of the Connection.
- Technology associated with the created Connection.



5.8 Create and Use an Oracle NetSuite Connection

Data Transforms uses the Oracle NetSuite JDBC Driver to connect to the Oracle NetSuite application. For Oracle NetSuite connections, Data Transforms allows you to load pre-built dataflows and workflows that you can run to transfer data from NetSuite to your target schema.

This topic has the following sections:

- Creating the Oracle NetSuite Connection
- Using the Build Data Warehouse Wizard
- Running the Pre-Built Workflows to Load Data into the Target Schema

Creating the Oracle NetSuite Connection

To define an Oracle NetSuite connection:

- 1. From the left pane of the Home page, click the **Connections** tab. **Connections** page appears.
- Click Create Connection. Create Connection page slides in.
- 3. Do one of the following:
 - In the Select Type field, enter the name or part of the name of the connection type.
 - Select the Applications tab.
- Select Oracle NetSuite as the connection type.
- Click Next.
- The Connection Name field is pre-populated with a default name. You can edit this value.
- In the JDBC URL textbox, enter the URL of the SuiteAnalytics Connect server to be used for the connection.
- 8. In the **User** text box enter the user name for connecting to the data server.
- 9. In the **Password** text box enter the password for connecting to the data server.
- 10. In the Account ID textbox, enter the account ID for connecting to the data server.
- 11. In the Role ID textbox, enter the role ID for connecting to the data server.
- 12. If the source that you want to use for mapping is a saved search, you need to also specify the following details in **Saved Search Extraction**:
 - Application ID: Enter the NetSuite Application ID for Data Transforms.
 - Version: Enter the NetSuite version number.
- 13. Select the checkbox in Build Data Model to install pre-built dataflows and workflows that you can run to extract data from NetSuite and move it to your Oracle target schema using the Build Data Warehouse wizard.
- **14.** Click **Test Connection**, to test the established connection.
- **15.** After providing all the required connection details, click **Create**. The new connection is created.



The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (**) next to the selected connection to perform the following operations:

- Select Edit, to edit the provided connection details.
- Select Test Connection, to test the created connection.
- Select Build Data Warehouse, to select the functional areas and create the NetSuite Data Warehouse in the target schema. See Using the Build Data Warehouse Wizard for more information.
- Select **Delete Schema**, to delete schemas.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- Name of the Connection.
- Technology associated with the created Connection.

Using the Build Data Warehouse Wizard

Data in your NetSuite account is grouped into business or subject areas in the Analytics Warehouse. The Build Data Warehouse wizard allows you to select the areas that you want to include in the newly created Data Warehouse.

To use the Build Data Warehouse Wizard:

- 1. On the Home page, click the **Connections** tab. The **Connections** page appears.
- Click the Actions icon (*) next to the Oracle NetSuite connection that you want to use to build the data warehouse and click Build Data Warehouse. The Build Data Warehouse wizard opens.
- **3.** From the **Connection** drop-down list, choose the Autonomous Database connection where your target schema resides.
- **4.** From the **Staging Schema** drop-down, all schema corresponding to the selected connection are listed in two groups:
 - Existing Schema (ones that you've imported into Oracle Data Transforms) and
 - New Database Schema (ones that you've not yet imported).

Select the schema that you want to use from the drop-down.

- Similarly select the Target Schema.
- 6. Click Next.
- Select the NetSuite Business Areas that you want to use to transfer data from the NetSuite Data Warehouse to the target schema.
- 8. Click Save.

Data Transforms starts the process to build the data warehouse. Click **Jobs** on the left pane of the Home page to monitor the progress of the process. When the job completes successfully, Data Transforms creates a Project folder that includes all the pre-built workflows and dataflows, which you can run to transfer data from the



NetSuite connection to your target schema. See Running the Pre-Built Workflows to Load Data into the Target Schema for more information.

Running the Pre-Built Workflows to Load Data into the Target Schema

When the Build Data Warehouse wizard completes successfully, Data Transforms creates a project that includes all the pre-built data flows and workflows that you can run to extract data from a Netsuite connection and load it into your target schema.

To view and run the pre-built workflows:

- Click Projects on the left pane of the Home page and select the newly created NetSuite project.
- Click Workflows in the left pane. The following pre-built workflows are listed in the Project Details page:
 - Stage NetSuite Source to SDS
 - Extract Transaction Primary Keys
 - Load SDS to Warehouse
 - Apply Deletes
 - All Workflows
- 3. Click the Actions icon (*) next to the workflow you want to run and click **Start**. Oracle recommends that you run **All Workflows** to execute all the pre-built workflows.

To see the status of the workflow, click **Jobs** from the left pane in the current project. When the job completes successfully, all the data from the NetSuite connection is loaded into the target schema.

5.9 Create an Oracle Object Storage Connection

You can use Data Transforms for uploading data from Oracle Object Storage to Autonomous Database.

To create an Oracle Object Storage connection you need to have an Oracle Cloud Infrastructure username and an auth token. See Getting an Auth Token for information about how to generate the auth token. You need to specify these details when you define the connection in Oracle Data Transforms.

Note the following:

- To use an Oracle Object Storage connection to import data into Data Transforms, you
 must use a public IP address to access the compute node. If you want to use a private IP
 address to access the Object Storage service, make sure that you have access to the
 Internet.
- The supported file format for loading data from Oracle Object Storage to Autonomous Database and vice versa is CSV.
- The supported data types are Numeric, Double, String, and Date.
- Data load is not supported for Oracle Object Storage connections.
- You cannot use an Oracle Object Storage connection as a target for mapping.

To define an Oracle Object Storage connection,



1. From the left pane of the Home page, click the **Connections** tab.

Connections page appears.

Click Create Connection.

Create Connection page slides in.

- 3. Do one of the following:
 - In the **Select Type** field, enter the name or part of the name of the connection type.
 - Select the **Databases** tab.
- 4. Select **Oracle Object Storage** as the connection type.
- 5. Click Next.
- **6.** The **Connection Name** field is pre-populated with a default name. You can edit this value.
- 7. Enter the URL in the **Object Storage URL** textbox. You can enter the URL in either of the following formats:
 - URL with fully qualified domain name. For example,

```
https://swiftobjectstorage.<your-region>.oraclecloud.com/v1/
<your-namespace>/<your-bucket>/<your-file>
```

```
https://objectstorage.<your-region>.oraclecloud.com/n/<your-
namespace>/b/<your-bucket>/o/<your-file>
```

 If you want to use the URL provided by the OCI Console, specify the URL only till the name of the bucket.
 For example,

```
https://swiftobjectstorage.<your-region>.oraclecloud.com/v1/
<your-namespace>/<your-bucket>
```

```
https://objectstorage.<your-region>.oraclecloud.com/n/<your-
namespace>/b/<your-bucket>/o
```

 If you choose Credential as the Connection Mode (see step 6), specify the URL till bucketname/o For example,

```
https://objectstorage./objectstorage./o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n//o/com/n/com/n//o/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/n/com/
```

The values for Region, Namespace and Bucket are auto-populated based on the URL provided.

- **8.** To select the **Connection Mode** do one of the following:
 - Select Swift Connectivity, and provide the following details:
 - In the User Name text box enter your Oracle Cloud Infrastructure username.



- In the **Token** text box enter the auth token.
- Select Credential and provide the ODI credential in the Enter Credential text box.
 You must create the credential in the repository and in the Autonomous Database
 that you created during instance creation. When you create a data flow to map data
 from Object Storage to Autonomous Database you need to create the ODI credential
 in target schema as well.

To create the credential, execute the following script:

9. Click Create.

The new connection is created.

10. Click **Test Connection**, to test the established connection.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (*) next to the selected connection to perform the following operations:

- Select Edit, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Select **Delete Schema**, to delete schemas.
- Select Delete Connection, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- Name of the Connection.
- Technology associated with the created Connection.

5.10 Create a REST Server Connection

You can connect to any REST service endpoint, import the data entities into Data Transforms, and use them as source in a data flow.

To create a generic REST connector, you need to provide the JDBC URL, username, and password to connect to the endpoint. You can also create and upload a config file that contains information such as the authentication methods, endpoints, and tables that you want to import data entities from.

The **Application** tab on the **Create Connection** page includes two connection options to create a generic REST connection - Generic REST and Generic REST Config. This topic has the following sections:

- Creating a Generic REST Connection
- Creating a Generic Rest Connection Using a Config File

The newly created connections are displayed in the **Connections** page.



Click the **Actions** icon (*) next to the selected connection to perform the following operations:

- Select Edit to edit the provided connection details.
- Select Test Connection to test the created connection.
- Select Delete Schema to delete schemas.
- Select Delete Connection to delete the created connection.

You can also search for the required connection to know its details based on the following filters:

- Name of the connection.
- Technology associated with the created connection.

Creating a Generic REST Connection

To create this connection you need to specify the REST service URL and choose a temporary schema where Data Transforms can create data entities after the reverse-engineering operation.

To define a REST server connection:

- 1. From the left pane of the Home page, click the **Connections** tab. **Connections** page appears.
- Click Create Connection.Create Connection page slides in.
- 3. Do one of the following:
 - In the **Select Type** field, enter the name or part of the name of the connection type.
 - Select the Applications tab.
- 4. Select **Generic Rest** as the connection type.
- 5. Click Next.
- The Connection Name field is pre-populated with a default name. You can edit this value.
- In the REST Service URL textbox, enter the URL of the endpoint that services the REST resources.
- **8.** In the **Proxy Host** textbox, enter the host name of the proxy server to be used for the connection.
- 9. In the **Proxy Port** textbox, enter the port number of the proxy server.
- 10. In the **User** text box enter the user name for connecting to the REST endpoint.
- 11. In the **Password** text box enter the password for connecting to the REST endpoint.
- **12.** Choose a connection from the **Staging Connection** drop-down list. The list displays only existing Autonomous Database connections. To use a different connection, create the connection before you reach this page.
- **13.** After providing all the required connection details, click **Test Connection** to test the connection.



14. Click Create.

The new connection is created.

Creating a Generic Rest Connection Using a Config File

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

To create a generic REST connector, you need the JDBC URL, username, password, and a config file. The config file is a model file with the <code>file_name.rest</code> naming convention that you need to upload when you create a REST Server connection. You need to specify the endpoints, table mappings, and the authentication methods to create the config file. You can create the config file using any text editor.

To define a REST server connection using a config file:

- From the left pane of the Home page, click the Connections tab.
 Connections page appears.
- Click Create Connection.Create Connection page slides in.
- 3. Do one of the following:
 - In the Select Type field, enter the name or part of the name of the connection type.
 - Select the Applications tab.
- 4. Select **Generic Rest Config** as the connection type.
- Click Next.
- **6.** The **Connection Name** field is pre-populated with a default name. You can edit this value.
- 7. Use the **Config File** text box to upload the config file that you want to use.
- 8. In the **JDBC URL** textbox, enter the URL to connect to the server.
- In the User and Password text boxes enter the user name and password for connecting to the REST endpoint. You may leave these fields blank if these values are not applicable or are already mentioned in the JDBC URL.
- **10.** After providing all the required connection details, click **Test Connection** to test the connection.
- 11. Click Create.

The new connection is created.



6

Data Loads

This section contains instructions for creating and running data loads.

Topics:

Create a Data Load

A data load allows you to load multiple data entities from a source connection to a target connection.

Run a Data Load

After you create the data load, you are taken to the Data Load Detail page that displays the details that you need to run a data load.

Monitor Status of Data Loads, Data Flows, and Workflows
 When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the
 background to complete the request. You can view the status of the job in the panel on
 the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor
 page.

6.1 Create a Data Load

A data load allows you to load multiple data entities from a source connection to a target connection.

You can create a data load in either of the following ways:

- Using the Create Data Load wizard.
- Using the Projects page.



Data load is not supported for Oracle Object Storage connections.

Creating a Data Load from the Home Page

This section describes the generic steps to create a data load. If you plan to load and transform data using OCI GoldenGate, create the data load using the Projects page. See Creating a Data Load from the Projects Page.

To create a data load from the Home page:

- From the left pane, click the Home tab. Click Load Data.
 The Create Data Load wizard appears.
- 2. In the Name field, enter a name for the data load. The field is pre-populated with a default name. You can edit this value.
- 3. Add a description. This is optional.

- 4. Select a project name from the drop-down. If this your first time here, click the + icon to create a project. If you have logged in as SUPERVISOR, the default project name is Home. For other users, the default project name is in the format <username>_Home. You can edit the default value. See Work with Projects for more information about projects.
- 5. Click Next.
- 6. To define your source connection, from the Connection drop-down, select the required connection from which you wish to add the data entities. Alternatively, click the + icon to create a new connection. See Work with Connections for more details about connections.
- 7. In the **Schema** drop-down, all schema corresponding to the selected connection are listed in two groups:
 - Existing Schema (ones that you've imported into Oracle Data Transforms) and
 - New Database Schema (ones that you've not yet imported).

Select the schema that you want to use from the drop down.



If there is missing information such as user name or password not specified, wallet missing, and so on, the list may fail to populate with a

"This connection has missing information." error. Click the Edit icon ($\[\]$) to open the Update Connection page where you can fill in the missing details.

- 8. Click Next.
- 9. Similarly, define the target connection.
- 10. Click Save.

The Data Load Detail page appears listing all the loaded data entities.

Creating a Data Load from the Projects Page

To create a data load from the Projects page,

- On the Home page, click **Projects**, and then the required project title. In the left pane, click **Data Loads**, and then click **Create Data Load**. The Create Data Load wizard appears.
- 2. In the Name field, enter a name for the data load.
- 3. Add a description. This is optional.
- Select the source and target schemas.



Make sure that you have created connections before you plan to create a data load using the Projects page. See Work with Connections for more details about connections.



- To use the OCI GoldenGate Deployment Console to load data entities:
 - Select the Use GoldenGate checkbox.
 The Create Data Load page now shows fields that are specific to OCI GoldenGate.
 - **b.** Select the GoldenGate connection from the drop-down.
 - To select the source connection, select the Registered Database and the Schema.
 - d. Similarly, define the target connection.
- To use all other connection types to load entities:
 - **a.** To define your source connection, select the connection from which you wish to add the data entities from the **Connection** drop-down.
 - **b.** From the **Schema** drop-down select the schema that you want to use. All schema corresponding to the selected connection are listed in two groups
 - Existing Schema (ones that you've imported into Oracle Data Transforms) and
 - New Database Schema (ones that you've not yet imported).
- 5. Click Create.

The Data Load Detail page appears listing all the loaded data entities.

6.2 Run a Data Load

After you create the data load, you are taken to the Data Load Detail page that displays the details that you need to run a data load.

It includes the details of the source schema, the data entities that are loaded from the source schema, and the details of the target schema. You can choose the action that you want to apply on each data entity – recreate, truncate, append - and load tables in bulk to the target schema. When you run a data load, multiple jobs run in the background to complete the request.

You can run a data load in either of the following ways:

- Using the Data Load Detail Page.
- Running a Data Load from the Workflow Details Page.



Data load is not supported for Oracle Object Storage connections.

Running a Data Load from the Data Load Detail Page

The Data Load Detail page displays the information that you need to run a data load. You can apply different actions - incremental merge, incremental append, recreate, truncate, append - on the data entities before loading it to the target schema.



Note:

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

If the data load is huge, you might want to increase the memory of the ODI Agent to avoid any issues. Follow the instructions in Increase the Memory of ODI Agent before you start to run the data load.

To run a data load from the Data Load Detail Page:

1. In the Data Load Detail page, select the data entities that you want to move to the target schema.

To filter the list, you can do one of the following:

- Enter the name or part of the name in the Name text box. This search returns
 data entities that include the specified string in the name. For example, the
 search string AD returns results such as ADMIN, ADDRESS, BADGE,
 UPGRADE, WORKLOAD, and so on.
- Turn on the Use Regular Expression toggle to use pattern matching to search for particular strings of characters. For example, the search string CO.* returns results such as CONTACT, COUNT, COUNTRY and so on.

You can select multiple data entities and load them to the target schema. You can also sort the displayed list using the following options:

- All Displays all the data entities that match the search criteria.
- Selected Displays all the rows that you selected.
- Unselected Displays all the unselected rows.
- Invalid Displays all the invalid rows.

Note:

These options display the list of data entities based on the search criteria. To view the list of all data entities, clear any applied filters.

- 2. Click on the required icon to choose any of the following actions:
 - Incremental Merge Updates the data in the selected column by comparing
 the source table with the target table based on the specified merge key. To use
 this option, select the column that you want to merge and then select the
 merge key. Click the Validate icon () to validate the selected values.
 - Incremental Append Updates data in the selected column in the target schema. To use this option, select the column that you want to update and click the Validate icon () to validate the selection.
 - Recreate If the table is already present in the target schema, drops the existing table and recreates it.



Note:

This option is not available for data entities that are loaded using OCI GoldenGate.

• Truncate – If the table is already present in the target schema, deletes all the data from the selected table. Nothing is dropped.

Note:

For Delta Share data loads the Data Load Detail page only includes the

- option. You cannot apply different actions incremental merge, incremental append, recreate, truncate, append on the data entities before loading it to the target schema. This is to make sure that the data is consistent between the Delta Sharing server and the target schema.
- Append If the table is already present in the target schema, adds rows to the table.
- Do Not Load Skips the selected data entity from the data load job. After you click
 Save, these data entities are no longer available for future data load jobs.

You can select multiple data entities and apply different actions. The unsaved rows are highlighted in bold.

Note:

These options are not available for Delta Share connections.

- 3. To specify how you want to store the source column names in the target tables, click Advanced Options, which is on the right side of the Data Load Detail page. Choose one of the following:
 - Retain original names by enclosing all names with delimiters Creates column names with the same names as is from the source tables in the target table.
 - Use no delimiters This is the default selection. Converts all the column names to upper case and replaces spaces and special characters with underscores.
 The following options are applicable to reserved words such as Date, Timestamp, Start, and so on.
 - Enclose with delimiters This is the default selection. Encloses column names that are reserved words with delimiters (not all column names).
 - Use a prefix Adds the specified prefix to column names that are reserved words (not all column names).

For column names that have the same name after conversion, the names are suffixed with a numeric value to maintain uniqueness. For example column names

```
Date, date, DATE, Item_@Code, Item$$Code, Item%$Code
```

are created in the target table as

```
DATE, DATE 0, DATE 1, ITEM CODE, ITEM CODE 0, ITEM CODE 1.
```



Note:

Once the data load is run, the selected options are applied and retained for all subsequent runs. You cannot change the configuration.

- Click to save the changes. A green checkmark (in the row indicates that the changes are saved.
- To start the data load,
 - Click .
 - For GoldenGate data loads, click

A confirmation prompt appears when the data load starts successfully.

To check the status of the data load, see the Status panel on the right below the Target Schema details. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows links to the jobs that execute to run this data load. Click the link to monitor the progress on the Job Details page. For more information about jobs, see Create and Manage Jobs.

Running a Data Load from the Workflow Details Page

You can add multiple data loads to a workflow along with data flows or workflows and run them as separate steps. The left panel of the Workflow Details page lists the data flows, workflows, and data loads that are available for use.

For data loads, the left panel lists the following two folders:

- Data Loads This folder lists all the data loads that you have created in the local Data Transforms instance. When you select a step in the workflow that is a local data load, the Properties Panel available on the right side of the design canvas displays the Type as Data Load and the Linked Object as Home><nameofDataLoad>.
- Remote Data Loads This folder lists all the data loads that you have created in a remote Data Transforms instance. See Create a Data Transforms Connection for Remote Data Load. When you select a step in the workflow that is a remote data load, the Properties Panel available on the right side of the design canvas displays the Type as Remote Data Load and the Linked Object as Name of the Data Transforms connection > <nameofDataLoad>.

After you have added the data loads to the workflow, click lacksquare to execute them.

6.3 Monitor Status of Data Loads, Data Flows, and Workflows

When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.



This panel includes the following:

- A Refresh icon ($\mathcal C$) to refresh the displayed status.
- The current status. You can see any of the following statuses:
 - Not Started If no data load, data flow, or workflow has been executed yet.
 - Running When the job starts.
 - Done When the job completes.
 The status of the last job run persists in this panel till you run a new data load, data flow, or workflow.



For data loads created using OCI GoldenGate, the status panel shows the link to the GoldenGate Deployment Console and the status of the Extract and Replicat processes.

• The name of the job and a link that takes you to the Jobs Details page where you can see detailed information about the job and monitor the execution status. For more information, see Create and Manage Jobs.

When you run a data load, multiple jobs run in the background to complete the request. This panel shows a link for each job that executes when running a data load.



7

Data Entities

Learn about how you can create and import data entities that you can then use to create data flows to transform data.

Topics:

Work with Data Entities

A Data Entity is a tabular representation of a data structure.

Import Data Entities

The easiest and most common way to create a Data Entity is by importing its structure from the connection type (for example, Oracle database).

Create Data Entities

You can manually create or update the Data Entities through the Oracle Data Transforms interface.

Create Data Entities within the Data Flow editor

If you have already created or imported your target data entity, then you would drag the data entity onto the Design Canvas and complete the column mappings and options.

View Statistics of Data Entities

The Preview tab displays detailed statistics of each data entity.

7.1 Work with Data Entities

A Data Entity is a tabular representation of a data structure.

It includes Database Tables or Views that can be used in a mapping as a source or target. They are simply the metadata for sources and targets. They are used in creating data flows.

You can add Data Entities to your newly created project in one of the following two ways:

- Importing Data Entities
- Creating Data Entities

All the newly created or imported Data Entities along with their details are displayed in the **Data Entities** page. The details include:

- Name of the Data Entity
- Connection for which the Data Entity was created
- Schema to which the Data Entity is associated
- Click the Actions icon (*) next to the selected Data Entity to perform the following operations:
 - Select Edit, to edit the existing details.
 - Select Preview, to preview the selected Data Entity. If the data entity belongs to an
 Oracle database, you can also view statistics of the table. See View Statistics of Data
 Entities for more details.

- Select **Delete**, to delete the selected Data Entity.
- To delete the Data Entities in bulk, in the Data Entities page, select the check boxes of the respective Data Entities and click **Delete**.
- You can also search for the required Data Entity to know its details based on the following filters:
 - Name of the Data Entity
 - Connection for which the Data Entity was created
 - Schema to which the Data Entity is associated
 - Tag that is associated with the Data Entity

7.2 Import Data Entities

The easiest and most common way to create a Data Entity is by importing its structure from the connection type (for example, Oracle database).

To import existing Data Entities:

- 1. From the left pane of the Home page, click the **Data Entities** tab. **Data Entities** page appears.
- 2. Click **Import Data Entities**, to import the existing data entities. **Import Data Entities** page slides-in.
- Select the Connection followed by Schema and then finally select the Type of Objects you want to import. For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.
- 4. [For Oracle Business Intelligence Cloud Connector (BICC) connections only] From the Offerings to import for collection, choose the offerings whose data stores you want to import. You must select at least one offering to import the BICC public view objects (PVO).



The import of BICC PVOs can take a long time depending on the number of selected objects. To improve performance, Oracle recommends that you use a mask to filter and limit the number of PVOs that you want to import.

Choose a Mask/filter if you don't want to import every object in the schema. Depending on the Connection Type, you will be presented with further options for importing.



For Oracle Object Storage connections, this value is case-sensitive. If **Batch similar files** is set to True, all the files that match the mask and have the same structure are grouped together into a single data entity.



- 6. [For Oracle Financials Cloud connections only] From the list in the Resources section, select the items that you want to import. When the import process completes, a table is created for each selected resource.
- 7. [For REST server connections only] In the **Resources** section, do the following:
 - In the Resource URI field, enter the URL of the REST service you want to import resources from.
 - Click the + icon.
 - In the Name column enter an identifier for the resource.
 - In the Operation URI column enter the URI of the resource.
 - Click Test Resource to check whether the entries are valid.

8. Click Start.

A Job is created and the corresponding Job ID is displayed for you to track the session. Click the Job ID to view the details of the job.

Upon successful execution of the job, all the selected Data Entities are imported. Click the Refresh icon $\mathcal C$ present at the right corner of the Data Entities page, to see the new imported Data Entities.

7.3 Create Data Entities

You can manually create or update the Data Entities through the Oracle Data Transforms interface.

Data entities should possess the corresponding objects in the source connection to be used as a source in a data flow. Usually the import process makes sure that these objects are in coordination. However, whenever you manually create or update Data Entities always make sure to check if both the definitions are in coordination with each other.

When you use a Data Entity as a target then it doesn't have to exist previously in the target connection and can be created as a part of Data Flow execution.

To create a new Data Entity:

- From the left pane of the Home page, click the **Data Entities** tab.
 Data Entities page appears.
- Click Create Data Entity, to create a new data entity.Create Data Entity page appears.
- 3. In the Name text box, enter the name of the new Data Entity that you are creating.
- **4.** From the **Connection** drop-down, select the required connection from which you wish to add the newly created Data Entity.



Oracle Financials Cloud connections are not listed here because you cannot manually create data entities for such connections. You can only import data entities from Oracle Financials Cloud REST endpoints using the Import Data Entities page. See Import Data Entities.



- 5. In the **Schema** drop-down, all schema corresponding to the selected connection are listed in two groups.
 - New Database Schema (ones that you've not imported from before) and
 - Existing Database Schema (ones that you've imported from before and are potentially replacing data entities).

From the **Schema** drop-down, select the required schema.



For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.

- **6.** From the **Type** drop-down, select the data entity type.
 - Table: To define the table structure for the newly created Data Entity, click the + icon to add columns. For each column, you can specify parameters such as Name, Data Type, Length, Scale, Not Null. Double click on the cell to configure the value.
 - Click the 'x' icon, to delete a row. Click the Up and Down arrows to sort the table rows.
 - Inline View: To create the data entity using inline code, enter the Select statement in the Query tab. For example, SELECT * FROM CUSTOMER. Click Validate.
 - The **Columns** tab displays a read-only list of the columns that the query returns. Click the **Preview** tab to see the column data.
- 7. In the **Tags** text box, enter a tag of your choice. You can use tags to filter the Data Entities displayed in the Data Entity Page.
- **8.** For Oracle Object Storage connections, this page displays the following options:
 - Contents Select the CSV file that contains the data you want to import. The
 metadata displayed in the preview table, such as the data type and length of
 columns, is based on the first row of the CSV file. Make sure that the CSV file
 has a header line. The header should contain only alphanumeric characters
 and no special characters.
 - Group Files Select this check box if you want to group data from multiple CSV files into one data entity. For example, say you want to merge data from Employee_Data1.csv, Employee_Data2.csv, and Employee_Data3.csv into a single data entity.
 - Resource Name Use this option along with Group Files. Enter the value you want to use to identify the files. The resource name should be a regular expression. You can use only an asterisk (*) as a wildcard character in the resource name. For example, Employee Data*.csv.
- For Oracle database connections, you can mark the data entities as a feature group. Expand Advanced Options and click the Treat as Feature Group checkbox.
- 10. Click Save.

The new Data Entity is created.



7.4 Create Data Entities within the Data Flow editor

If you have already created or imported your target data entity, then you would drag the data entity onto the Design Canvas and complete the column mappings and options.

To create the definition of a new target table while in the Data Flow editor,

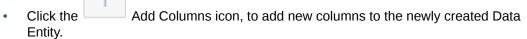
- 1. Select the component at the end of your data flow.
- 2. Click the Add Data Entity icon present on the top right corner of the target component.
- 3. Add Data Entity page appears allowing you to configure the following details of the target component:

General tab

- In the Name text box, enter the name of the newly created Data Entity.
- The Alias text box is auto-populated with the name of the newly created Data Entity.
- From the **Connection Type** drop-down, select the required connection from which you wish to add the newly created Data Entity.
- It loads the server name coined at the time of connection creation. From the Server drop-down, select the required server name from which you wish to add the newly created Data Entity.
- From the Schema drop-down, select the required schema.
- Click Next.

Columns tab

It allows you to create, remove or edit the column definitions.



A new column is added to the displayed table.

- The table displays the following columns:
 - Name
 - Data Type Click the cell to configure the required Data Type.
 - Scale
 - Length
 - Actions Click the cross icon to delete the created column.
- To delete the columns in bulk, select the columns and click the Delete icor
- To search for the required column details, in the Search text box enter the required column name and click enter. The details of the required column are displayed.
- Click Next.

Preview Data Entity tab



It displays a preview of all the created columns and their configured details. If the data entity belongs to an Oracle database, you can also view statistics of the table. See View Statistics of Data Entities for more information.

Click Save.

The new target Data Entity is created.

- **5.** Expand the **Properties Panel** in the right pane to view the following settings of the created components:
 - General Displays the Name of the component along with its Connection and Schema details.
 - Attributes Displays the details of all the attributes associated to the component.
 - Column Mapping Click Auto Map to map all the columns automatically.
 - Preview Click to have a preview of the component.
 - Options Change the options as appropriate.

7.5 View Statistics of Data Entities

The Preview tab displays detailed statistics of each data entity.

Note:

This feature is available for Oracle database tables only.

You can view the statistics of a selected data entity in one of following ways:

- In the Data Entities list, click the **Actions** icon (*) next to the Data Entity and click **Preview**. Select the **Statistics** tab to view the statistics of the selected data entity.
- On any data flow click on any source or target data entity, and expand the properties panel in the right pane. Click Preview.

The statistical data is presented as follows:

- The total number of rows and columns in the data entity is displayed at the top.
- The statistics panel displays the thumbnail graphs for each column with information about the Min, Max, Distinct, and Null values.
- Two types of thumbnail representations are displayed based on the histogram:
 - A bar chart represents data for frequency and top-frequency histograms. The bar chart show the first top 10 values for the number of rows in the table.
 - A table lists data for Hybrid and Height-Balanced histograms. The table displays the entire data and is scrollable. The table displays the range for the values and the percentage of rows in each range.
- You can click each thumbnail to view the statistics of the column in a new browser tab.
- The detailed view of each chart also shows the type of histogram.



8

Projects

The section includes information about managing and working with projects.

Topics:

Work with Projects

A project is the top-level container, which can include multiple folders to organize your data flows or work flows into logical groups.

8.1 Work with Projects

A project is the top-level container, which can include multiple folders to organize your data flows or work flows into logical groups.

The Projects page displays all the existing projects, which includes projects created by all users not just the logged in user.

You can perform the following operations of a Project folder:

- View Details
- · Edit the name
- Delete

Creating a Project

To create a new project:

- 1. Click Create Project on the Projects page.
- 2. On the Create Project page, provide a project name.
- 3. Click Create.

After creating the project, you're redirected to the **Projects** page.

Viewing Project Details

To view project details:

- On the Projects page, select a project and click View Details from the project's Actions icon (*) or select a project to open its details page.
- It displays all the resources associated to the selected project. It includes details such as Data Flow Name and Folder details for the selected project.
- 3. Click the **Actions** icon () of the respective data flow to edit, rename, copy, change folder, start or delete the required data flow of the selected project.



Deleting a Project

When you delete a project, you also delete the resources it contains. Once you delete a project, it cannot be restored.



Be sure to review all contained resources before you delete the project.

To delete a project, on the Projects page, select Delete from the **Actions** icon (*) for the project you want to delete or select a project to open its details page, and then click **Delete**.

In the **Delete Project** dialog, click **Delete** to confirm the delete operation.



9

Data Flows

This section contains detailed information about creating and working with data flows, which includes adding components, mapping columns, and executing a data flow.

Topics:

About Data Flow Editor

The Data flow editor is divided into five parts, the Data Entity Panel, the Transformations Toolbar, the Design Canvas, the Properties Panel, and the Status Panel.

Create a Data Flow

A data flow defines how the data is moved and transformed between different systems.

Add Components

Add the data entities and database functions to the Design Canvas, and connect them in a logical order to complete your data flows.

Component Properties

The Properties Panel displays various settings for components selected in the Design Canvas.

Map Data Columns

When you connect the source data entity with the target data entity, the column names are automatically mapped by the column names. You have a choice to map the columns by Position or by Name or map the columns manually using the Expression Editor.

Validate and Execute a Data Flow

After your mappings are ready, you can proceed to validate and execute the data flow.

Schedule Data Flows or Workflows

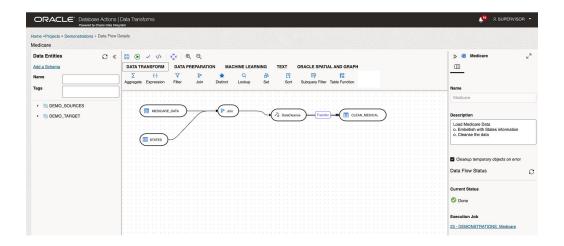
You can schedule workflows and data flows to run at specified time interval.

• Monitor Status of Data Loads, Data Flows, and Workflows

When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.

9.1 About Data Flow Editor

The Data flow editor is divided into five parts, the Data Entity Panel, the Transformations Toolbar, the Design Canvas, the Properties Panel, and the Status Panel.



- Data Entities Panel: The data entity panel displays the Data Entities that are
 available to use in your Data flows. The displayed list can be filtered using the
 Name and Tags fields. The panel includes options that let you add schemas,
 import data entities, remove any of the schemas that are associated with the data
 flow, and refresh data entities. See Add Components for information about how to
 use these options.
- Database Functions Toolbar: The Database Functions toolbar display the
 database functions that can be used in your data flows. Just like Data Entities, you
 can drag and drop the Database tools you want to use on the design canvas. See
 Supported Database Functions for more information.
- **Design Canvas**: The design canvas is where you build your transformation logic. After adding the Data Entities and Database Functions to the design canvas, you can connect them in a logical order to complete your data flows.
- Properties Panel: The properties panel displays the properties of the selected object on the design canvas. The Properties Panel is grouped into four Tabs.
 General, Attributes, Preview Data, Column Mapping, and Options. Not all tabs are available as they vary based on the selected object. See Component Properties to know more about these options.
- Status Panel: When you run a data flow, the Status Panel shows the status of the job that is running in the background to complete the request. You can see the status of the job that is currently running or the status of the last job. For more information about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows.

After designing the required data flow,

- Click $\stackrel{\square}{=}$, to save the created/designed data flow.
- Click $\stackrel{\triangleright}{\searrow}$, to align the nodes of the designed data flow.
- Click lacksquare , to execute the created data flow.
- Click $\stackrel{\bigoplus}{}$, to maximize or minimize the created data flow diagram in the design canvas.



9.2 Create a Data Flow

A data flow defines how the data is moved and transformed between different systems.

A data flow in Data Transforms connects sources to targets through a flow of components such as Join, Filter, Aggregate, Set, Split, and so on. See Supported Database Functions for more information.

When you run a data flow, Data Transforms uses the joins, filters, mappings, and constraints to transform source data and load it to target tables. Note that you can run only one execution flow at a time. You cannot put multiple flows on a Data Flow and a flow cannot diverge into multiple flows.

You can create data flows in any of the following ways:

- From the Projects page
- From the Data Flows page within a project
- From the Home page

From the Projects page

To create a data flow from the Projects page,

- On the Projects page, click Create Data Flow.
 Create Data Flow page appears:
- 2. In the Name field, enter a name for the new data flow.
- Select Create New Project, if you wish to create new project folder for the newly created data flow.
- Else, click Add to Existing Projects, if you wish to add the newly created data flow to an existing project folder.
- If you have selected Create New Project for the previous option, in the Project Name field, enter the name of the newly created project.
- Else, if you have selected Add to Existing Projects for the previous option, select the required project from the Project Name drop-down arrow.
- 7. In the **Description** field, enter a description for the newly created data flow.
- 8. Click Create.

From the Data Flows page within a project

To create a data flow from the Data Flows page within a project,

- On the Projects page click the project tile you wish to create a new data flow for. The Project Details page appears.
- In the Data Flows page, click Create Data Flow.
- 3. Provide the **Name** and **Description** of the new data flow.
- 4. Click Next.
- **5.** To define your source connection, from the **Connection** drop-down, select the required connection from which you wish to add the data entities.
- 6. In the Schema drop-down, all schema corresponding to the selected connection are listed in two groups:



- Existing Schema (ones that you've imported into Oracle Data Transforms) and
- New Database Schema (ones that you've not yet imported).

Select the schema that you want to use from the drop down. For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.

7. Click Save.

The Data Flow Editor appears that allows you to create a new data flow.

From the Home page

To create a data flow from the Home page,

- 1. On the Home page, click **Transform Data**. The Create Data Flow page appears.
- 2. Provide the Name and Description of the new data flow.
- 3. Select a project name from the drop-down. Alternatively, click the + icon to create a project.
- 4. Click Next.
- 5. From the **Connection** drop-down, select the required connection from which you wish to add the data entities. Alternatively, click the + icon to create a new connection.
- 6. In the Schema drop-down, all schema corresponding to the selected connection are listed in two groups:
 - Existing Schema (ones that you've imported into Oracle Data Transforms) and
 - New Database Schema (ones that you've not yet imported).

Select the schema that you want to use from the drop down.

7. Click Save.

9.3 Add Components

Add the data entities and database functions to the Design Canvas, and connect them in a logical order to complete your data flows.

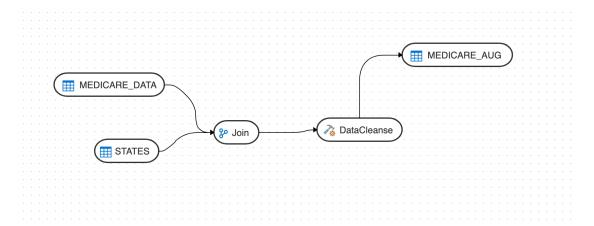
To add components to your data flow:

- 1. In the Data Entities panel, click **Add a Schema** to add schemas that contain the data entities that you want to use in the data flow.
- 2. In the Add a Schema page, select the connection and schema name.
- 3. Click Import.
- 4. In the Import Data Entities page, select the Type of Objects you want to import. Choose a Mask/filter if you don't want to import every object in the schema and click Start.
- 5. The Data Entities panel lists the imported data entities. The panel includes various options that let you do the following:
 - Refresh Data Entities Click the Refresh icon $\mathcal C$ to refresh the displayed list.
 - Name Search for data entities by name.
 - Tags Filter the data entities by the name of the tag used.



- Import Data Entities Right-click the schema to see this option. Use this option to import the data entities.
- Remove Schema Right-click the data entity to see this option. Use this option to remove the schema from the list. Note that this option does not delete the schema, it only removes the association of the schema with this data flow.
- 6. Similarly add more schemas to the Data Flow, if required.
- 7. Drag the required Data Entities that you want to use in the data flow and drop them on the design canvas.
- **8.** From the **Transformations** toolbar, drag the transformation component that you want to use in the data flow and drop them on the design canvas.
- Select an object on the design canvas, and drag the Connector icon () next to it to connect the components.
- 10. After saving the data flow, there may be a Transfer icon overlaid on one or more of the component connections. This indicates that ODI has detected an additional step and it is required to move the data between data servers. You can click on this Icon to view properties associated with this step.

For example:



9.4 Component Properties

The Properties Panel displays various settings for components selected in the Design Canvas.

Depending on the component selected, you may see any of the following icons:

- General () Displays the name of the component along with its connection and schema details. You can edit some of these properties.
- Attributes ($\stackrel{\blacksquare}{\blacksquare}$) Displays the details of all the attributes associated with the component.
- Column Mapping () Allows you to map all the columns automatically. See Map Data Columns for more information.



- Preview () Displays a preview of the component. For Oracle tables, you can also view the statistics of the selected data entity. See View Statistics of Data Entities for details about the statistical information available.
- Options () Displays options such as
 - Truncate Table Replaces any existing target table content with new data.
 - Append Inserts records from the flow into the target. Existing records are not updated.
 - Incremental Integrates data in the target table by comparing the records of the flow with existing records and updating the records when their associated data is not the same. Those that don't yet exist in the target are inserted. The option includes an Auto compression feature that is set to True by default. For data flow jobs that use the Incremental Update mode to load data onto a compressed Oracle target partition, the Auto compression feature recompresses the modified target partitions after the load completes successfully. For table partitions that are not originally compressed, the compression is skipped irrespective of whether Auto compression is set to true.

Note:

The Auto compression option is available to the ADMIN user or to a user with the DWROLE role. For data flows that have schema users other than ADMIN you need to either assign the DWROLE to the user or disable Auto compression to avoid execution errors.

9.5 Map Data Columns

When you connect the source data entity with the target data entity, the column names are automatically mapped by the column names. You have a choice to map the columns by Position or by Name or map the columns manually using the Expression Editor.

To map columns by Position or by Name:

- 1. Select the target Data Entity.
- 2. Click the arrow icon present on the top right corner to expand the Properties Panel. This will give you more space to work with.
- 3. In the Properties Panel, click the Column Mapping icon ().
- **4.** To map the columns by Position or by Name, from the **Auto Map** drop-down menu, select **By Position** or **By Name**.

To map the columns manually:

- 1. From the **Auto Map** drop-down menu, select **Clear** to clear the existing mappings.
- Drag and drop the attributes from the tree on the left to map with the Expression column.



3. To edit an expression, click the **Edit** icon of the respective column. The **Expression Editor** appears allowing you to perform the required changes (for example, you can just add an expression-"UPPER" or open the Expression Editor to edit the expression).



Use the expression editor only if you have complex expressions for a particular column.

Click OK.

9.6 Validate and Execute a Data Flow

After your mappings are ready, you can proceed to validate and execute the data flow.

Do the following:

- Click Save.
 After saving, if data needs to be staged before transforming, Transfer button is added to one or more links. You can click these buttons to set more options, if available.
- 2. Click the Simulate Code icon () if you want to check the code that will run to complete the tasks that are performed when you execute the data flow job. The source and target details are displayed in different colors for ease of reference. This is handy if you want to check if the mapping is correct before you run the job or if the job fails. Note that the code cannot be used for debugging. For detailed information about the job, see the Job Details page.
- 3. Click the **Validate** icon (✓) in the toolbar above the design canvas to validate the data flow.
- 4. After a successful validation, click the **Execute** icon () next to the **Validate** icon to execute the data flow.

A message appears that displays the execution Job ID and name. To check the status of the data flow, see the Status panel on the right below the Properties Panel. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel also shows the link to the Job ID that you can click to monitor the progress on the Jobs page. For more information, see Create and Manage Jobs.

For data flows created using Oracle Object Storage connections, the data from the source CSV file is loaded into the target Oracle Autonomous Database. You can also export data from an Oracle Autonomous Database table to a CSV file in Oracle Object Storage.

9.7 Schedule Data Flows or Workflows

You can schedule workflows and data flows to run at specified time interval.

Each schedule allows you to specify a start date and time, and execution frequency.

To schedule a data flow or a workflow:

- 1. In the left pane, click **Schedules**.
- 2. Click Create Schedule.



- **3.** From the **Resource** drop-down menu, select **Data Flow** or **Workflow** as appropriate.
- **4.** From the **Resource Name** drop-down list, select a data flow or workflow that you want to schedule.
- 5. From the **Frequency** drop-down list, select the time frame in which you wish to execute the created schedule. You can schedule it On Startup, Simple, Daily, Hourly, Weekly, Monthly or Yearly. Based on the selected frequency, **Time** field appears allowing you to schedule the selected Data Flow or Workflow. For detailed information refer to the below table:

Time Values	
NIL	
Click Select Date Time icon next to Date and Time field, to select the required date and time in MM/DD/YY and HH:mm format.	
Click the clock icon next to Time field, to select the required time in HH:mm format.	
Click up and down arrows next to Time field, to select the required time in Minutes: Seconds format.	
 Run Every parameter appears with check-boxes for all the days of the week. Click the clock icon next to Time field, to select the required time in HH: mm format. 	
 Select the required date from the Monthly Date drop-down box. Click the clock icon next to Time field, to select the required time in HH: mm format. 	
 From the Monthly date drop down box select the required value. From the Week day drop-down box, select the required weekday in which you wish to schedule the job. 	
From the Month drop down box, select the required month in which you wish to schedule the job. From the Month by data drop down box.	
 From the Monthly date drop down box, select the required date in which you wish to schedule the job. Click the clock icon next to Time field, to select the required time in HH: mm 	

- 6. In the Number of Attempts on Failure text box, enter the required value or click the up or down arrows next to the text box, to select the required value. This value denotes the number of retry attempts that should happen after the schedule failure.
- Stop Execution After field denotes the time after which the schedule has to stop after executing it. It can be in Hours, Minutes or Seconds. Select the value and its unit.



- 8. Select the status of the scheduled Data Flow or Workflow from the **Status** options. It can be **Active**, **Inactive** or **Active for Period**. Through **Active for Period** option you can configure the exact time frame of the created schedule. It's Starting and Ending Date with time, time interval in which you wish to run the schedule and exceptions, if any.
- **9.** After configuring the above details, click **Save**.

The created schedules get listed in the **Schedules** page along with the following details:

- Resource
- Status
- Frequency
- Valid or Invalid

Click the Actions menu () of the respective schedule to perform the following operations:

- Click Edit to edit the details of the created schedule.
- Click **Disable** to disable the created schedule. The status of a schedule becomes
 inactive when you disable it and gets displayed in the Scheduled list. Select the
 schedule and click **Enable**, to enable it whenever required. You can enable it again
 for a specific period, through **Enable for Period** option.
- Click Validate to validate the created schedule.
- Click **Delete** to delete the created schedule. Upon confirmation the created schedule is deleted.

To schedule a data flow or a workflow specific to a project:

- Click the Project title displayed on the Projects page.
 Project Details page appears.
- In the left pane, click Schedules.
 Schedules page appears. It displays all the schedules pertaining to the selected project.



This page is project specific and displays only the Data Flows and Workflows pertaining to the selected project.

- To create a new schedule, click Create Schedule.
- Enter all the required details and click **Save**.

 The new schedule is created and is added to the existing list of schedules for the project.

9.8 Monitor Status of Data Loads, Data Flows, and Workflows

When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.

This panel includes the following:

- ullet A Refresh icon (${\cal O}$) to refresh the displayed status.
- The current status. You can see any of the following statuses:



- Not Started If no data load, data flow, or workflow has been executed yet.
- Running When the job starts.
- Done When the job completes.
 The status of the last job run persists in this panel till you run a new data load, data flow, or workflow.

Note:

load.

For data loads created using OCI GoldenGate, the status panel shows the link to the GoldenGate Deployment Console and the status of the Extract and Replicat processes.

The name of the job and a link that takes you to the Jobs Details page where you
can see detailed information about the job and monitor the execution status. For
more information, see Create and Manage Jobs.
 When you run a data load, multiple jobs run in the background to complete the
request. This panel shows a link for each job that executes when running a data



10

Workflows

This section contains detailed information about creating and working with workflows.

Topics:

Introduction to Workflows

A workflow is made up of multiple flows organized in a sequence in which they must be executed.

Create a New Workflow

You can add data flows, workflows, or data loads in a workflow.

Schedule Data Flows or Workflows

You can schedule workflows and data flows to run at specified time interval.

Monitor Status of Data Loads, Data Flows, and Workflows

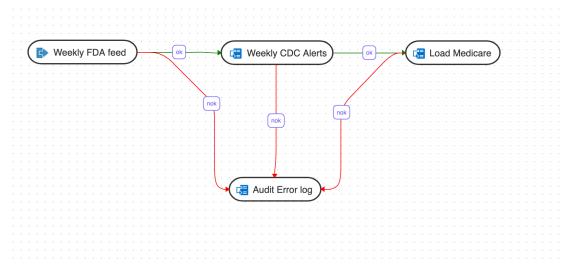
When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.

10.1 Introduction to Workflows

A workflow is made up of multiple flows organized in a sequence in which they must be executed.

Each data flow is executed as a step. You can also add workflows as well as SQL queries as steps within a workflow. When you execute a workflow, a data flow either succeeds or fails. Depending on whether the first data flow succeeds or fails, you can choose the next data flow that must be executed.

Here is an example of a workflow:



In this example, the workflow performs the following actions:

- 1. Execute the "Weekly FDA feed" data flow.
- If the "Weekly FDA feed" data flow execution is successful, execute the "Weekly CDC alerts" data flow.
- If the "Weekly CDC alerts" data flow execution is successful, execute the "Load Medicare" data flow.
- 4. If any of the above data flow fails, then execute the "Audit error log" data flow.

10.2 Create a New Workflow

You can add data flows, workflows, or data loads in a workflow.

To create a new workflow for your project:

- 1. On the **Home** page, click the required Project title. You are navigated to the **Project Details** page.
- 2. In the left pane, click **Workflows**. The Workflow page appears.
- 3. On the Workflow page, click Create Workflow. The Create Workflow page appears.
- In the Name field provide a name for the new workflow and click Create.
 The new workflow is created and listed in the Workflow page with its Name and Folder details.
- 5. Click the Workflow to configure the Workflow Details.
- 6. From the left panel drag the data flows, workflows, or data loads that you want to run in the workflow. If you have connected to any other Data Transforms instance, you can also add data loads that you have created in that Data Transforms instance. See Create a Data Transforms Connection for Remote Data Load and Run a Data Load for more information.
- Select either the ok (green arrow) icon, the not ok (red arrow) or the ok/not ok (black arrow) in the toolbar.

This defines the mode that subsequent links drawn on the canvas will be set to.

- ok (green) defines the success path.
- not ok (red) defines the failure path.
- ok/not ok (black) defines a path that is followed on success or failure.

Use the Sleep (clock icon) to add a delay in the workflow. Drag the Sleep icon on the canvas and connect it in the flow with either the ok (green), not ok (red) or ok/not ok (black) links. This will add a delay at that point in the flow.

8. If you wan to add a SQL or PL/SQL query as a step in the workflow, do the following:



- Drag the SQL icon (^{SQL}) on the canvas.
- Double click the SQL step in the editor to open the step properties page.
- Select the Attributes tab.
- From the **Connection** drop-down select the connection you want to run the query on.





The drop-down lists only Oracle database connections.

- In the SQL textbox add the query that you want to run.
- Select the step and click the Connector icon () next to it to connect it with the next step.
- 10. After defining all the required Workflow details,
 - Click , to save the created/designed workflow.
 - Click K, to align the nodes of the designed workflow.
 - Click , to execute the created workflow.
 - Select a single step in the canvas and click the Execute Step icon (), to execute only the selected data flow or workflow.
 To check the status of the workflow, see the Status panel on the right below the

Properties Panel. For details about the Status panel on the right below the Properties Panel. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows the link to the Job ID that you can click to monitor the execution status on the Jobs page. For more information about jobs see Create and Manage Jobs.

• Click , to maximize or minimize the created workflow diagram in the design canvas.

The newly created workflows get listed in the Project Details page. The following details are displayed:

- Name of the Workflow
- Folder corresponding to the workflow Default Folder

To view the details of a workflow, click the name of the workflow and you are navigated to the Workflow Details page.

In the left pane, you can search for the required Data Flow or Workflow using the Name filter. In the **Name** text box, enter the name of the required Data Flow or Workflow.

Select a step in the canvas and check the **Properties Panel** available at the right side of the design canvas to know the following properties of the selected step in a created Data Flow or Workflow:

- Name
- Linked Object



You cannot edit this field.

Step -



 First Step - Select the First Step check-box, to execute the selected step as the first step for execution in a Data Flow or Workflow.



You can select only a single step as a first step for execution in a Data Flow or Workflow.

- Number of attempts on Failure
- Time between attempts in seconds(s)
- Log steps in journal Select any of the following options from the drop-down box:
 - Always
 - Never
 - Error

10.3 Schedule Data Flows or Workflows

You can schedule workflows and data flows to run at specified time interval.

Each schedule allows you to specify a start date and time, and execution frequency.

To schedule a data flow or a workflow:

- 1. In the left pane, click Schedules.
- 2. Click Create Schedule.
- **3.** From the **Resource** drop-down menu, select **Data Flow** or **Workflow** as appropriate.
- From the Resource Name drop-down list, select a data flow or workflow that you want to schedule.
- 5. From the Frequency drop-down list, select the time frame in which you wish to execute the created schedule. You can schedule it On Startup, Simple, Daily, Hourly, Weekly, Monthly or Yearly. Based on the selected frequency, Time field appears allowing you to schedule the selected Data Flow or Workflow. For detailed information refer to the below table:

Frequency	Time Values
On Startup	NIL
Simple	Click Select Date Time icon next to Date and Time field, to select the required date and time in MM/DD/YY and HH:mm format.
Daily	Click the clock icon next to Time field, to select the required time in HH:mm format.
Hourly	Click up and down arrows next to Time field, to select the required time in Minutes: Seconds format.



Frequency	Time Values
Weekly	 Run Every parameter appears with check-boxes for all the days of the week. Click the clock icon next to Time field, to select the required time in HH: mm format.
Monthly (day of the month)	 Select the required date from the Monthly Date drop-down box. Click the clock icon next to Time field, to select the required time in HH: mm format.
Monthly (week day)	 From the Monthly date drop down box select the required value. From the Week day drop-down box, select the required weekday in which you wish to schedule the job.
Yearly	 From the Month drop down box, select the required month in which you wish to schedule the job. From the Monthly date drop down box, select the required date in which you wish to schedule the job. Click the clock icon next to Time field, to select the required time in HH: mm

- 6. In the **Number of Attempts on Failure** text box, enter the required value or click the up or down arrows next to the text box, to select the required value. This value denotes the number of retry attempts that should happen after the schedule failure.
- 7. **Stop Execution After** field denotes the time after which the schedule has to stop after executing it. It can be in **Hours**, **Minutes** or **Seconds**. Select the value and its unit.
- 8. Select the status of the scheduled Data Flow or Workflow from the **Status** options. It can be **Active**, **Inactive** or **Active for Period**. Through **Active for Period** option you can configure the exact time frame of the created schedule. It's Starting and Ending Date with time, time interval in which you wish to run the schedule and exceptions, if any.
- After configuring the above details, click Save.
 The created schedules get listed in the Schedules page along with the following details:
 - Resource
 - Status
 - Frequency
 - Valid or Invalid

Click the Actions menu () of the respective schedule to perform the following operations:

- Click Edit to edit the details of the created schedule.
- Click **Disable** to disable the created schedule. The status of a schedule becomes inactive when you disable it and gets displayed in the Scheduled list. Select the schedule and click **Enable**, to enable it whenever required. You can enable it again for a specific period, through **Enable for Period** option.



- Click Validate to validate the created schedule.
- Click **Delete** to delete the created schedule. Upon confirmation the created schedule is deleted.

To schedule a data flow or a workflow specific to a project:

- Click the Project title displayed on the Projects page.
 Project Details page appears.
- In the left pane, click Schedules.
 Schedules page appears. It displays all the schedules pertaining to the selected project.



This page is project specific and displays only the Data Flows and Workflows pertaining to the selected project.

- To create a new schedule, click Create Schedule.
- Enter all the required details and click Save.
 The new schedule is created and is added to the existing list of schedules for the project.

10.4 Monitor Status of Data Loads, Data Flows, and Workflows

When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.

This panel includes the following:

- A Refresh icon (\mathcal{C}) to refresh the displayed status.
- The current status. You can see any of the following statuses:
 - Not Started If no data load, data flow, or workflow has been executed yet.
 - Running When the job starts.
 - Done When the job completes.
 The status of the last job run persists in this panel till you run a new data load, data flow, or workflow.



For data loads created using OCI GoldenGate, the status panel shows the link to the GoldenGate Deployment Console and the status of the Extract and Replicat processes.



- The name of the job and a link that takes you to the Jobs Details page where you can see detailed information about the job and monitor the execution status. For more information, see Create and Manage Jobs.
 - When you run a data load, multiple jobs run in the background to complete the request. This panel shows a link for each job that executes when running a data load.



11

Jobs

This section includes information about creating and working with jobs.

Topics:

Create and Manage Jobs

When you execute a data load, data flow or workflow Oracle Data Transforms creates a job to complete the process in the background.

11.1 Create and Manage Jobs

When you execute a data load, data flow or workflow Oracle Data Transforms creates a job to complete the process in the background.

A job is made up of multiple steps that corresponds to an execution task.

The **Jobs** page by default, lists the execution job sessions that are running and completed for the present day as Date parameter is set to Today by default. You can click on any of the instance id from the list to see its details.

Create a Job from the Home page

To create a new job,

- On the Home page, click **Jobs** on the left pane to access the Jobs page.
- Click Create Job.
 The Create Job page slides-in.
- From the Resource Type drop-down, select the type of resource for which you wish to create a new job. It can be - Data Flow, Workflow or Schema.
- All the resources associated with the selected Resource Type get listed in the Resource field. From the Resource drop-down, select the required resource for which you wish to create a new job.
- Click Create.

A new job is created and is added to the existing list of jobs in the Jobs page.

For each session, you can see the Job Session ID, Name, Status, Start Time and End Time for each of the jobs that are executed.

Create a Job from the Projects page

To view the job sessions pertaining to a specific project:

- Click the Project tile displayed in the Projects page.
 Project Details page appears.
- In the left pane, click Jobs.
 Jobs page appears. It displays all the jobs and their details pertaining to the selected project.

Note:

This page is project specific and displays only the jobs pertaining to the selected project.

- To create a new job, click Create Job.
- Enter all the required details and click Create.
 A new job is created and is added to the existing list of jobs for the project.

View and Delete a Job

Click the Actions menu () to **View Details** of the jobs, **Delete** and **Rerun** the jobs, if required. If any of the steps failed, you can click on them to see the error details.

To delete a job, select the check box of the respective job session and click **Delete**. Upon confirmation the selected job gets deleted.

Search for a Job

You can also search for the required job session to know its details based on filters such as:

- Name Name of the Job
- **Status** Select the required session status from the list All, Done, Error, Queued, Running, Waiting and Warning.
- Date Select the date in which the required job session was executed All, Today, Yesterday, Last Week, Last Month and Custom Range, which allows you to select specific From and To dates.



12

Users

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

This section provides information for creating and managing users.

Topics:

- Create Users
- Change User Passwords
- Delete Users

12.1 Create Users

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

The User Management page of the Administration tab of Oracle Data Transforms helps you to manage users.

A user corresponds to the login name used to connect to a repository, and can have privileges depending on the role assigned. The SUPERVISOR user is assigned the Oracle Data Integrator Admin (odiadmin) role by default. The odiadmin role allows users to:

- Assign the odiadmin role to a user
- Change user password
- Delete users

The SUPERVISOR user can assign this role when creating new users.

To create a new user:

- In the left pane, click Administration.
 A warning message appears.
- 2. Click Continue.
- In the left pane, click User Management.User Management screen appears, which displays the list of users.
- 4. Click Create User.
- 5. In the Create User page, enter the name. Maximum length is 50 characters.
- Select the Allow managing other users checkbox to allow this user to reset passwords or delete users.



Note:

This option is only available if you are the SUPERVISOR user or have been assigned the odiadmin role.

7. Enter and reenter and password for the new user.

Note:

The password should:

- contain alpha-numeric characters,
- be 6-12 characters in length,
- not start with a number,
- contain at least one alphabet in upper case,
- contain at least one special character. Only \$, # and are allowed.

8. Click Create User.

The newly created user appears in the list.

12.2 Change User Passwords

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

You must have the odiadmin role assigned to be able to change the password of any of the listed users. If you are not assigned the odiadmin role, you can only change the password of the logged in user.

In this case, the **Actions** menu is available only for the logged in user. You cannot change the password of any of the other users, even the ones that you have created.

To change the user password:

- On the User Management page, click Change Password from the user's Actions icon (*).
- 2. Enter the new password and click **Update User**.

To change the password of the SUPERVISOR user:

- On the User Management page, click Change Password from the user's Actions icon (*).
- Enter the new password and click Update User. You are logged out of Data Transforms.
- 3. Restart the jetty server. Log in as OPC user and execute the following commands:

```
ssh -i <path to id_rsa> opc@<Instance IP>
sudo su
```



```
systemctl stop|start jettyserver.service
exit
```

4. Launch Oracle Data Transforms from the Database Actions page.

12.3 Delete Users

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

You must have the odiadmin role assigned to be able to delete any of the listed users (except the SUPERVISOR).

To delete a user, you need to assign ownership of the objects created by this user to any of the other listed users.



You cannot delete users that have jobs in the running state.

If you are not assigned the odiadmin role, the **Delete** option is not available in the Actions menu.

To delete a user:

- 1. On the User Management page, click **Delete** from the user's Actions icon ().
- 2. Select a new owner for the objects created by this user from the **Other Users** drop-down.
- 3. Click Delete.



Autonomous Database

This section provides details about accessing Oracle Data Transforms from Oracle Autonomous Database.

Topics:

- Manually Register the Created Oracle Data Transforms Instance to Autonomous Database
- Map Autonomous Database User to Data Transforms User
- Unregister the Oracle Data Transforms Instance from Autonomous Database
- Access Oracle Data Transforms From Database Actions page
- Access Oracle Data Transforms From Data Studio
- Enable Access to Private Data Sources from Autonomous Database

13.1 Manually Register the Created Oracle Data Transforms Instance to Autonomous Database

APPLIES TO: • Data Transforms instance that is registered with Autonomous Database.

The RegisterODI.sh script allows you register your Oracle Data Transforms instance with Autonomous Database.

To manually register the Oracle Data Transforms instance to Autonomous Database after creating your instance,

- Connect to the compute instance running Oracle Data Transforms as described in Connect to Oracle Data Transforms Instance.
- Execute the RegisterODI.sh file using the following commands:

```
cd u01/oracle/transforms_home/common/scripts
sh RegisterODI.sh
```

- A terminal window appears allowing you to enter the following details of the Autonomous Database instance which you had already created in Create the Oracle Data Transforms Instance step:
 - Admin username User name of the Autonomous Database instance.
 - Admin password Password of the Autonomous Database instance.
 - URL Scheme Press Enter to select the default value http.
 - IP Address IP address of the computer which you wish to register to Autonomous Database.





You are prompted to provide all the above details repeatedly until you provide a valid IP address.

- Oracle Data Transforms Port It is the port in which Oracle Data Transforms runs on the local computer. Press Enter to select the default value - 9999.
- Map the database user (created during ODI provisioning) to Oracle Data
 Transforms User. See Map Autonomous Database User to Data Transforms User for details.

After configuring the above details, the created instance is successfully registered to the Autonomous Database. After successful registration, you can launch Oracle Data Transforms directly from the Autonomous Database, Database Actions page. For more details on this, refer to Access Oracle Data Transforms From Database Actions page.

In some cases, you may encounter an error ("401 Not Authorised") when launching Oracle Data Transforms from the Database Actions page. To resolve this issue, you will need to restart the Jetty server.

To stop, start, or restart the jetty server, login as OPC user and execute the following commands:

```
ssh -i <path to id_rsa> opc@<Instance IP>
sudo su
systemctl stop|start jettyserver.service
exit
```

13.2 Map Autonomous Database User to Data Transforms User

APPLIES TO: • Data Transforms instance that is registered with Autonomous Database.

The Administration tab of Oracle Data Transforms helps you to map the database user (created during provisioning) to Oracle Data Transforms User.

The administrator in Oracle Data Transforms is called the SUPERVISOR. By mapping the users, you can directly access the Oracle Data Transforms application from Database Actions page. For more details, refer to Access Oracle Data Transforms From Database Actions page.

For mapping the users,

- In the left pane, click Administration. A warning message appears.
- Click Continue.User Mapping screen appears.
- 3. In the Database URL text box, enter the Database URL created during Autonomous Database instance creation, for example - https:// u6dx8gx07phfkil-db202003100909.adb.us-phoenix-1.oraclecloudapps.com/. To get this URL,



- From OCI console, click the Hamburger icon and navigate to Autonomous
 Database and select the autonomous database instance for which you created the repository.
- Click the Service Console and navigate to Development → SQL Developer. You
 can find the database URL here and copy it to the Database URL text box.
- 4. Click the Add icon () to add a new column in the User Mapping table.
- Double click the **Database User** column to create a user mapping for the required Database User.
- 6. Double click the **ODI User** column and click the drop-down arrow to list all the ODI users available for mapping in Oracle Data Transforms application.
- 7. From the **ODI User** drop-down select the required ODI user, to whom you wish to map the newly added Database user.
- 8. Click **Save**.

 A confirmation message regarding the user mapping update appears.
- 9. Click the Refresh icon () next to the + icon, to see the newly created user mapping column in the User Mapping table.
- 10. If you wish to delete the created user mapping, select the user and click the Action icon(**) next to the user and select **Delete**.
- For deleting user mappings in bulk, select the respective user check-boxes and click Delete.

Upon confirmation the selected user mappings are deleted.

You can search for the required user mapping from the User Mapping table, through Name filter.

Enter the name of the user mapping that you are looking for in the Name text box.
 The respective user mapping column is displayed.

13.3 Unregister the Oracle Data Transforms Instance from Autonomous Database

APPLIES TO: • Data Transforms instance that is registered with Autonomous Database.

The UnregisterODI.sh script allows you unregister your Oracle Data Transforms instance from Autonomous Database.

To unregister your Oracle Data Transforms instance from Autonomous Database,

- Connect to the compute instance running Oracle Data Transforms as described in Connect to Oracle Data Transforms Instance.
- Navigate to your Data Transforms desktop and double click the Unregister with ADP shortcut icon.
 or



Execute the UnregisterODI.sh file using the following commands:

cd u01/oracle/transforms_home/common/scripts
sh UnregisterODI.sh

- A terminal window appears allowing you to enter the following details of the Autonomous Database instance which you had already created in Create the Oracle Data Transforms Instance step:
 - Admin username User name of the Autonomous Database instance.
 - Admin password Password of the Autonomous Database instance.

After providing the above details, the instance is successfully unregistered from the Autonomous Database. The Oracle Data Transforms will no longer be available in the Autonomous Database, Database Actions page.

13.4 Access Oracle Data Transforms From Database Actions page

APPLIES TO: • Data Transforms instance that is registered with Autonomous Database.

The Autonomous Database users can access the Oracle Data Transforms interface from the Data Tools tab present in Database Actions page, if the DB hosts the Oracle Data Transforms repository.

Follow the below steps to access Oracle Data Transforms:

- 1. On the OCI console, click the Hamburger icon present on the top left corner.
- Click Autonomous Data Warehouse.You are navigated to Autonomous Databases in ODI Compartment page.
- Select the Autonomous Database used for creating the ODI repository (for example, ODIQAADW1).
 - You are navigated to the respective Autonomous Database Details page.
- **4.** From the Autonomous Database Details page, click the **Service Console** tab. You are navigated to the Autonomous Data Warehouse page.
- 5. From the left pane of the Autonomous Data Warehouse page, click **Development**.
- In this page, select SQL Developer Web.The Oracle Database Actions login page appears.
- 7. In the Oracle Database Actions login page, in the respective Username and Password fields enter the login credentials of the Autonomous database user who is linked to the Oracle Data Transforms instance created earlier.
- After entering the login credentials, click Sign In. The Oracle Database Actions page appears.
- On the Oracle Database Actions page, click the Hamburger icon present on the top left corner.
- 10. From the left pane click Data Tools and select Data Transforms. The Oracle Data Transforms login page appears. Upon successful login, Oracle Data Transforms page appears allowing you to perform all the data transformations.



13.5 Access Oracle Data Transforms From Data Studio

APPLIES TO: • Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

Data Transforms combines all the elements of data integration - data movement, data synchronization, data quality, and data management to ensure that information is timely, accurate, and consistent across complex systems.

To use the Data Transforms tool you must access Database Actions as the ADMIN user or have the DATA_TRANSFORM_USER role assigned. See Manage User Profiles with Autonomous Database for information on granting roles.

To access the Data Transforms tool:

- Login as an ADMIN user or a user with the DATA_TRANSFORM_USER role assigned.
- Click Data Transforms in the Database Actions page, or click the Selector icon and select Data Transforms from the Data Studio menu in the navigation pane.



Data Transforms is also available as a separate listing on OCI Marketplace. If you have already registered a Data Transforms instance from OCI Marketplace with the Autonomous Database, the Data Transforms card on the Database Actions page will continue to take you to your Marketplace instance. If you wish to use the embedded Data Transforms, then you must unregister the Marketplace instance. See Unregister the ODI Instance from Autonomous Database.

- 3. When you login to the Data Transforms tool for the first time, you need to provide the database user credentials to sign in. It takes approximately 1-2 minutes for the Data Transforms service to start. After the service starts, the Data Transforms home page opens.
- 4. From the left pane of the Home page, click the **Connections** tab to view the newly created Autonomous Database connection.
 - 4. Click the Actions icon next to the connection and select Edit.
- In the Update Connection page, enter the database username and password to use the connection. Click Update to save the changes.

13.6 Enable Access to Private Data Sources from Autonomous Database

APPLIES TO: • Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

The Autonomous Database from where you are accessing Data Transforms should be configured to use a private endpoint to be able to communicate with private database



sources. Otherwise when you try to create and test such a connection, you might get a "failed to connect" error.

If your Autonomous Database is configured to use a Private Endpoint, then you can only access private data sources from clients in the same Virtual Cloud Network (VCN). See Configuring Network Access with Private Endpoints for detailed instructions.

After you have configured network access using private endpoints, you need to do the following additional configurations for Oracle Data Transforms to communicate with private data sources. This chapter includes the following topics:

- Configuring DNS and Traffic Management
- Configuring RCE Proxy on the Private Endpoint

Configuring DNS and Traffic Management

You need to set up a fully qualified domain name (FQDN) and specify this FQDN in the connection host field when you create the connection. Refer to https://docs.oracle.com/en-us/iaas/Content/DNS/Concepts/views.htm for more information.

The VCN DNS resolver should know about the FQDN and be able to resolve it to a Type A IP address.

If you are not able to test a connection in Data Transforms on an Autonomous Database that is configured to use private endpoints, then do the following:

- Check whether the query to set the ROUTE_OUTBOUND_CONNECTIONS database property to PRIVATE_ENDPOINT is run. See Enhanced Security for Outbound Connections with Private Endpoints for more information.
- Make sure that the FQDN is used as the host so that the framework resolves the FQDN to the required IP address.
 Ensure that the data source VM has the FQDN set up with the required security rules.

If test connection fails even after using an FQDN, then do the following:

- From the OCI services menu, click DNS Management under Networking.
 Configure private DNS zones views and resolvers.
- Create and publish a record using the FQDN and the private IP (Type A) of the target database.



Do not use *.oraclecloud.com as the domain name when you set up the FQDN because it is a reserved domain name.

Retry test connection.
 The test connection operation should complete successfully.

Configuring RCE Proxy on the Private Endpoint

For Autonomous Databases that are configured to use a private endpoint, the Autonomous Database service supports RCE to access resources in the customer subnet through a private route.



If you encounter a "failed to connect" error when you try to test a connection, contact your database administrator to check whether the database is RCE-enabled.

Note that RCE proxy can forward the reverse traffic only to the default OCI DNS resolver. If there is any overridden resolver, then connectivity will fail.



Reference

This section includes information about Oracle Data Tranforms such as the installation location, log file information as well as instructions for resetting the schema user password, using Run Commands in your Data Transforms instance, and increasing the memory of the ODI Agent.

Topics:

Terminology Information

There are similarities in the functionality between Oracle Data Integrator and Oracle Data Transforms, however, the terminology is different.

Installation Location

You may need this information if you want to change any aspect of the installation.

Log Files Location

This topic lists the log files that help you to keep a track record on all the events happening in your Oracle Data Transforms instance on Oracle Cloud Marketplace.

Reset the Schema User Password

The schema password expires every 180 days from the repository creation date or from the last password reset date. When you try to connect to such an existing repository the Data Transforms instance provisioning fails. You need to reset the password in the Autonomous Database as well the local wallet.

- Use the OCI Console to Run Commands in the Data Transforms Instance
 When the Oracle Data Transforms repository provisioning process completes, certain commands are available for use in the compute instance Run command feature.
- Increase the Memory of ODI Agent
- Backup and Restore Data Transforms Repositories
 You can move repositories from one database instance to another or backup and restore
 your existing repositories. Note that this is applicable only for Oracle repositories.

14.1 Terminology Information

There are similarities in the functionality between Oracle Data Integrator and Oracle Data Transforms, however, the terminology is different.

The differences in the terminology used in Oracle Data Integrator Studio and Oracle Data Transforms are:

Oracle Data Integrator Studio	Oracle Data Transforms	Functions
Data Server	Connection	Represents the physical object where data is stored.
		Examples include Database Instances, File Servers, and Cloud Application Instances. The properties contain all information required to connect and access data.
Technology	Connection Type	Examples include Oracle, IBM DB2, Oracle Netsuite, Oracle Object Storage.
		For the list of supported technologies, see Supported Connection Types.
Physical Schema	Schema	Examples include Database Schema, File server Folder, Object Storage Bucket.
Data Store	Data Entity	A Tabular representation of a data structure. Examples include Database Tables, Files.
Reverse Engineering	Import Data Entity	The process of obtaining the metadata for a filtered set of objects in a Schema. Examples include Table Definitions from an Oracle Schema, VOs from a Cloud Applications Offering.
Model	Data Entity	The container for the imported objects. The properties define the rules for importing (how to filter and so on).
Project	Project	Container for the Transformation Design components (Mappings, Packages, Jobs). This doesn't include Data Servers, Schemas, Data Stores and Folders (which are shared between Projects)
Mapping	Data Flow	The Transformation Design. Defines how data flows from one Data Store to another, and how it is transformed.
Package	Workflow	Defines the sequence in which Mappings will be executed, together with what happens of failure conditions.
Job / Session	Job	The execution of a Package, Mapping or Reverse Engineer.



Oracle Data Integrator Studio	Oracle Data Transforms	Functions
Schedule	Schedule	The rules for when a session executes.
Resource	Resource	Something you can schedule (a mapping or package).

14.2 Installation Location

You may need this information if you want to change any aspect of the installation.

The following installation location is used by the Oracle Data Transforms image.

Table 14-1 Installation Location

Area	Location on Server
MW_HOME	/u01/oracle/transforms_home

14.3 Log Files Location

This topic lists the log files that help you to keep a track record on all the events happening in your Oracle Data Transforms instance on Oracle Cloud Marketplace.

Log Files	Location
Agent application Logs	/u01/oracle/transforms_home/logs/app.log
Rest application Logs	/u01/oracle/transforms_home/logs/ odi_adp_rest_txt.log
Instance provisioning Logs (used only during provisioning)	/u01/oracle/transforms_home/logs/ odi_lean_install.log
Repo creation Logs (used only during provisioning)	/u01/oracle/transforms_home/logs/repoCreation.log
Configuration Logs (used only during provisioning)	/u01/oracle/transforms_home/logs/ transformsStartup.log

14.4 Reset the Schema User Password

The schema password expires every 180 days from the repository creation date or from the last password reset date. When you try to connect to such an existing repository the Data Transforms instance provisioning fails. You need to reset the password in the Autonomous Database as well the local wallet.

To reset the schema user password:

1. Login to the Autonomous Database that was used for creating the ODI repository.



2. Run the following command to reset the password for the user:

alter user username identified by newpassword

- 3. Connect to the compute instance that is running Oracle Data Transforms.
- 4. Navigate to the Data Transforms desktop and double click the updatepassword.desktop shortcut icon. In the terminal window that appears, enter the new password.

This updates the password in the local wallet and restarts the jetty server.



Skip steps 3 and 4 if the instance that you want to connect to is no longer available and you are creating a new instance using the existing repository.

14.5 Use the OCI Console to Run Commands in the Data Transforms Instance

When the Oracle Data Transforms repository provisioning process completes, certain commands are available for use in the compute instance Run command feature.

You can use these commands as templates to create new commands that you can run without having to login into the instance.

Name	Command	Description
restartCommand	<pre>sudo systemctl restart jettyserver.service</pre>	Restarts the jetty server.
archiveCommand	<pre>cp -R /u01/oracle/ transforms_home/logs /tmp</pre>	Archives the log files to the local temporary folder.
listLogs	ls /u01/oracle/ transforms_home/logs	Lists the log files.
displayStartupLo g	<pre>cat /u01/oracle/ transforms_home/logs/ transformsStartup.log</pre>	Displays the configuration logs. This information is useful for troubleshooting any configuration issues.

To use the run command feature, you have to setup the following policy:

Allow dynamic-group odi_group to manage instance-family in compartment odi
Allow dynamic-group odi_group to use instance-agent-command-family in compartment odi
Allow dynamic-group odi_group to use instance-agent-command-execution-family in compartment odi where request.instance.id=target.instance.id

To use the pre-installed scripts to create commands to run on an instance:

 On the OCI console, open the navigation menu present on the top left corner and click Compute.



- 2. Under Compute, click Instances.
- 3. Click the instance that you want to run the command on.
- Under Resources, click Run command. The Run command screen lists the commands that you can use as a template to create and execute commands.
- 5. Find the command in the list, click the Actions menu, and then click **View command** details.
- **6.** On the Command details page, copy the command from the **Command Content** section and close the page.
- 7. Click Create command.
- 8. Enter a name for the command.
- 9. In the Add script section, select Paste script and paste the copied command in the box.
- 10. In the Output type section, select the location to save the output of the command.
- 11. Click Create command.
- **12.** Wait for the completion of the command. See the **Execution Status** column to monitor whether the command has run successfully.

If the command output was saved to an Object Storage location, either download the response object from the bucket where it was saved or navigate to the Object Storage preauthenticated request URL. If the command output was saved as a plain text file, find the command in the list, click the Actions menu, and then click **View command details** to view the output.

14.6 Increase the Memory of ODI Agent

APPLIES TO: • Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

If the data that you are loading from the source schema is huge, then you may want to increase the memory of the ODI Agent to avoid OutOfMemory exception errors.

To increase the memory of the ODI Agent:

- Edit the /u01/oracle/transforms_home/common/scripts/jettyServer.sh file.
- 2. Add the java -Xms1024m -Xmx4096m parameter.
- 3. Restart the jetty server. Log in as OPC user and execute the following commands:

```
ssh -i <path to id_rsa> opc@<Instance IP>
sudo su
systemctl stop|start jettyserver.service
exit
```

14.7 Backup and Restore Data Transforms Repositories

You can move repositories from one database instance to another or backup and restore your existing repositories. Note that this is applicable only for Oracle repositories.

The topics within this section describe how you can:



- Use SQL Developer to Export and Import Data Transforms Repositories
 Use the SQL Developer Data Pump feature to export the source schema into a dump file and then import the dump file into another repository.
- Switch to the Imported Repository
 To switch to the imported repository, you need to update the odi-setup.properties and repository.properties files with information related to the imported repository.
- Update the Imported Repository using DataTransformsPostImport.sh
 Use the DataTransformsPostImport.sh utility to make the necessary updates to the imported repository.

14.7.1 Use SQL Developer to Export and Import Data Transforms Repositories

Use the SQL Developer Data Pump feature to export the source schema into a dump file and then import the dump file into another repository.

This topic has the following sections:

- Export Source Schema
- Copy the Dump File to Oracle Cloud Object Storage
- Create the Credential Object to Connect to Object Storage
- Create the Pre-Authenticated URL
- Import the Data Transforms Repository dump
- Stop the jetty server

Follow the below steps to import the ODI repository dump:

Export Source Schema

Run the Export Wizard from SQL Developer to export the source schema into a dump file, which you will later import into the target repository.

To export the Data Transforms repository:

- Start SQL Developer.
- 2. Define a connection to connect to the source repository. Set up proxy, if needed.
- 3. Use this connection to open the DBA panel.
- 4. Add the newly created connection.
- Expand the Data Pump node, right-click Export Jobs and select Data Pump Export Wizard. This opens the Export Wizard.
- **6.** In the **Source** screen, provide the following information:
 - Data or DDL: Select Data and DDL from the drop-down list.
 - Types to be Exported in Data Pump: Select Schemas.
- Click Next.
- In the Schemas screen select the source schema that you want to export. Click Next.



- 9. Leave the default settings in the Filter and Table Data screens as is and click Next.
- 10. In the Options screen select the Enable Logging checkbox to enable logging. This generates an export log file, which will be saved in the DATA_PUMP_DIR directory. You can either leave the default log file name as is or change it. Click Next.
- 11. In the **Output Options** screen, the Choose Output Files section displays a list of output files with the .DMP extension. Select the file that you want to export. This dump file will also be saved in the DATA_PUMP_DIR directory. Select the **Append Timestamp to Dump, Log and Job names** radio button and click **Next**.
- 12. In the **Job Schedule** screen, specify information about the job such as the Job Name and Job Description. Select the required scheduling options or accept the default settings to run the job immediately. Click **Next**.
- **13.** In the **Summary** screen, review the displayed information.
- 14. Click Finish to start the export process.
 A new export job will be added in the DBA window. Select the job from the Export Jobs node to check the status. You can also see the execution log in the lower-right window.

When the job completes successfully, the Export Log and the dump file can be found in the **DATA_PUMP_DIR** directory.

Copy the Dump File to Oracle Cloud Object Storage

You need to copy the exported dump file to Oracle Cloud Object Storage so that you can access it from the target repository. You can use the Object Storage of either the source or target tenant. This will allow you to import the .dmp file from any tenant, as long as you have the corresponding authentication token in order to access its Object Storage. To get the exported file, use the PL/SQL PUT OBJECT function.

- Open a SQL worksheet connected to the source database where the data pump export happened.
- Run the following function in SQL Developer:

```
BEGIN
    DBMS_CLOUD.PUT_OBJECT(credential_name => 'ODI',
        object_uri => ' https://objectstorage.<your-
region>.oraclecloud.com/n/<your-namespace>/b/<your-bucket>/o/
SourceSchemaName.dmp ',
        directory_name => 'DATA_PUMP_DIR',
        file_name =>'NameOfDumpFile.DMP');
END:
```

Here

- The credential name "ODI" should have already been created in the source database. Check all the existing credentials from USER CREDENTIALS.
- The value for object_uri can be obtained from the properties of the target bucket on Object Storage.

Create the Credential Object to Connect to Object Storage

Create a credential object on the target repository to access the Object Storage to import the dump file. The first step is to create an authentication token in Oracle Object Storage. See Getting an Auth Token for information about how to generate the auth token. Copy and save

the generated token. You will use this token to create the credential object to connect to the Object Store.

To create the credential object, execute the following script:

```
BEGIN

DBMS_CLOUD.CREATE_CREDENTIAL(

    credential_name => 'Credential name',
    username => 'oci username',
    password => '<Enter the Authentication token here>');
END;
```

Create the Pre-Authenticated URL

To access the DMP files from Object Storage you need to create pre-authenticated requests for each file that will be used for the import. When you create a Pre-authenticated Request for an object in Oracle Object Storage, you get the URL to the object. Copy and save this URL. You need to provide this URL in the Data Pump Import Wizard to import the dump file. See Creating a Pre-Authenticated Request for information about how to create the URL.

Import the Data Transforms Repository dump

Run the Import Wizard from SQL Developer to import the dump file into the target repository.

To import the repository dump:

- 1. Start SQL Developer.
- 2. Define a connection to connect to the target repository.
- 3. Use this connection to open the DBA panel.
- 4. Add the newly created connection.
- Expand the Data Pump node, right-click Import Jobs and select Data Pump Import Wizard. This opens the Import Wizard.
- **6.** In the **Type** screen, provide the following information:
 - Data or DDL: Select Data and DDL from the drop-down list.
 - Type of Import: Select Schemas.
 - In Choose Input Files set the location of the DMP file(s) by selecting the Credential that you created in Create the Credential Object to Connect to Object Storage and provide the pre-authenticated URI that you created in Create the Pre-Authenticated URL.
- 7. If the credential connection is successful, the input file is scanned, and the Filter screen displays all the available DMP schemas. Select the source schema(s) that you would like to import into the target repository. Click Next.
- 8. In the **Remapping** screen map the Source and Destination schemas. Click **Add Row**. The source schema will be automatically displayed. Enter the destination schema. The source schema will be imported to this empty schema. Click **Next**.
- In the Options screen select the Enable Logging checkbox to enable logging. This generates an import log file, which will be saved in the DATA_PUMP_DIR



- directory. You can either leave the default log file name as is or change it. Click Next.
- 10. In the Schedule screen, specify information about the job such as the Job Name and Job Description. Select the required scheduling options or accept the default settings to run the job immediately. Click Next.
- **11.** In the **Summary** screen, review the displayed information.
- 12. Click Finish to start the import process. A new import job will be added in the DBA window. Select the job from the Import Jobs node to check the status. You can also see the execution log in the lower-right window.

When the job completes successfully, the Import Log and the dump file can be found in the **DATA PUMP DIR** directory.

Stop the jetty server

After you have imported the repository, you must stop the jetty server. To stop the jetty server, login as OPC user and execute the following commands:

```
ssh -i <path to id_rsa> opc@<Instance IP>
sudo su
systemctl stop jettyserver.service
exit
```

14.7.2 Switch to the Imported Repository

To switch to the imported repository, you need to update the odi-setup.properties and repository.properties files with information related to the imported repository.

To switch to the imported repository:

1. Create the odi-setup.properties file in /u01/oracle/transforms_home/common/scripts and add the following properties. If the file already exists, replace the existing content of the file with these details.

```
dbTech=ADB
rcuCreationMode=false
odiSchemaPassword=<valid password>
odiSchemaUser=<odi schema username>
odiSupervisorPassword=<odi SUPERVISOR password>
walletZipLoc=<path_to_zipped_wallet>
workRepoName=<WORK REPO NAME>
adwInstancePassword= <adw Instance password>
```

Here,

- workRepoName=<WORK REPO NAME> is an optional property but you may have to configure this property if your default work repository name is not WORKREP.
- adwInstancePassword= <adw Instance password> is an optional property but configure this property only when you have used OPTACH for applying a patch on your ODI instance and wish to run Upgrade Assistant (UA) using the configuration script odiMPConfiguration.py.



2. Create the repository.properties file in /u01/oracle/transforms_home/common/scripts and add the following properties. If the file already exists, replace the existing content of the file with these details.

```
masterReposDriver=oracle.jdbc.OracleDriver
masterReposUrl=jdbc\:oracle\:thin\:@<new work repository name>_high?
TNS_ADMIN\=/u01/oracle/transforms_home/wallets/wallet_<wallet file
name>
workReposName=<WORK REPO NAME>
masterReposUser=<odi schema username>
```

14.7.3 Update the Imported Repository using DataTransformsPostImport.sh

Use the DataTransformsPostImport.sh utility to make the necessary updates to the imported repository.

Before you run the utility make sure that the repository properties file has the correct information that points to the imported master repository. See Switch to the Imported Repository for more information.

The DataTransformsPostImport.sh file is located in /u01/oracle/transforms_home/common/scripts. If the script is not found, create the file using the following properties:

```
#!/bin/sh
export install_home=/u01/oracle/transforms_home

CLASSPATH=.:$install_home/apps/*:$install_home/deploy/jettydeploy/
APPADMINREST/webapp/WEB-INF/lib/ojdbc8-19.3.0.0.jar:$install_home/
deploy/jettydeploy/APPADMINREST/webapp/WEB-INF/lib/
oraclepki-19.3.0.0.jar:$install_home/deploy/jettydeploy/APPADMINREST/
webapp/WEB-INF/lib/osdt_cert-19.3.0.0.jar:$install_home/deploy/
jettydeploy/APPADMINREST/webapp/WEB-INF/lib/osdt_core-19.3.0.0.jar
export CLASSPATH

/u01/oracle/jdk1.8.0_211/bin/java -cp $CLASSPATH
oracle.odi.setup.util.mp.DataTransformsPostImport

exec $SHELL
```

After you have successfully run the utility, you must start the jetty server. To start the jetty server, login as OPC user and execute the following commands:

```
ssh -i <path to id_rsa> opc@<Instance IP>
sudo su
systemctl start jettyserver.service
exit
```

Connect to Data Transforms and verify whether the repository has been successfully imported.

