

Database Actions

Using Guide for On-Premises Oracle Database



Release 26.2

G58172-03

August 2026

ORACLE®

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>.

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Contents

Preface

Audience	i
Product Accessibility	i
Related Documents	i
Conventions	i
Third-Party License Information	ii

1 Changes in Release 26.2 for Using Oracle Database Actions

2 About Oracle Database Actions

2.1 About the Database Actions User Interface	1
2.1.1 Using the Omniseach Bar	3
2.2 Accessing Database Actions	5
2.3 Signing In to Database Actions	6
2.4 The Home Page - Launchpad	7

Part I Development

3 The SQL and Data Modeler Pages

3.1 The Overview Page	1
3.2 The SQL Page	2
3.2.1 Object Navigator and Files	2
3.2.2 Executing SQL Statements in the Code Editor	6
3.2.2.1 About Session State in Database Actions	9
3.2.2.2 Keyboard Shortcuts	9
3.2.2.3 Support for Multilingual Engine	11
3.2.3 Viewing the SQL Output	13
3.2.3.1 Using the Explain Plan Diagram	16
3.2.4 Loading Data	19
3.2.4.1 Loading Data from a Local File to a New Table	19

3.2.4.2	Loading Data from a Local File to an Existing Table	21
3.2.4.3	Loading Data from Multiple Local Files into Multiple Tables	24
3.3	The Data Modeler Page	26
3.3.1	Navigating Diagrams and Objects	27
3.3.2	About the Data Modeler Editor	29
3.3.2.1	About Viewing JSON Data Guides	31
3.4	Creating and Editing Database Objects	35
3.4.1	The Table Properties Dialog	35
3.4.1.1	Columns Pane	37
3.4.1.2	Primary Key Pane	40
3.4.1.3	Unique Keys Pane	40
3.4.1.4	Indexes Pane	41
3.4.1.5	Foreign Keys Pane	41
3.4.1.6	Table Constraints Pane	42
3.4.1.7	Comments Pane	42
3.4.1.8	Storage Pane	42
3.4.1.9	External Table Properties Pane	42
3.4.1.10	Materialized View Pane	45
3.4.1.11	DDL Pane	47
3.4.1.12	Output Pane	47
3.4.2	The Index Properties Dialog	47
3.4.3	The Sequence Properties Dialog	48
3.4.4	The View Properties Dialog	49
3.4.5	The Synonym Properties Dialog	50
3.4.6	The Materialized View Log Properties Dialog	51
3.4.7	Implied Foreign Keys	53

4 The REST Pages

4.1	About RESTful Web Services	1
4.2	RESTful Services Terminology	1
4.3	About the Overview Page	2
4.4	About the AutoREST Page	3
4.4.1	Enabling REST Access for a Database Object	4
4.4.2	Disabling REST Access for a Database Object	4
4.4.3	Generating cURL Requests for a REST-Enabled Database Object	5
4.5	About the REST Search Toolbar	5
4.6	Creating RESTful Web Services	5
4.6.1	Managing Resource Modules	6
4.6.1.1	Creating a Resource Module	7
4.6.1.2	Importing a Resource Module	8
4.6.1.3	Editing a Resource Module	8

4.6.1.4	Deleting a Resource Module	8
4.6.1.5	Publishing/Unpublishing a Resource Module	8
4.6.1.6	Exporting a Resource Module	9
4.6.1.7	Viewing the Module in OpenAPI View	9
4.6.2	Managing Resource Templates	9
4.6.2.1	Creating a Resource Template	10
4.6.2.2	Editing a Resource Template	10
4.6.2.3	Deleting a Resource Template	11
4.6.3	Managing Resource Handlers	11
4.6.3.1	Creating a Resource Handler	12
4.6.3.2	Editing a Resource Handler	13
4.6.3.3	Deleting a Resource Handler	13
4.6.3.4	Viewing Resource Handler Details and Managing Parameters	13
4.6.3.5	Using Pre-Authenticated Requests to Access ORDS RESTful Services	18
4.6.4	Example: Inserting a Record using a POST Handler	19
4.7	Securing RESTful Web Services	21
4.7.1	Managing Roles	21
4.7.1.1	Creating a Role	22
4.7.1.2	Editing a Role	22
4.7.1.3	Deleting a Role	22
4.7.1.4	Viewing Assigned Privileges	22
4.7.2	Managing Privileges	22
4.7.2.1	Creating a Privilege	23
4.7.2.2	Editing a Privilege	23
4.7.2.3	Deleting a Privilege	24
4.7.3	Managing OAuth Clients	24
4.7.3.1	Creating an OAuth Client	25
4.7.3.2	Editing an OAuth Client	26
4.7.3.3	Deleting an OAuth Client	26
4.7.3.4	Exporting an OAuth Client	26
4.7.3.5	Managing Secrets	26
4.7.4	Examples	27
4.7.4.1	Creating an OAuth Client Using the Client Credentials Grant Type	27
4.7.4.2	Creating an OAuth Client Using the Auth Code Grant Type	28

5 The Liquibase Page

5.1	Generate a Deployment	1
5.2	About the Liquibase User Interface	2

6 The Charts and Dashboards Page

6.1	Overview	1
6.2	Creating or Editing a Chart	2
6.3	Example: Creating and Editing a Chart	3
6.4	Creating or Editing a Dashboard	6

7 The JSON Page

7.1	About the JSON User Interface	1
7.2	Managing JSON Collections	4
7.2.1	Creating a Collection	4
7.2.2	About Adding or Editing a JSON Document	5
7.2.2.1	About Database Differences in JSON Documents	5
7.3	About Querying Documents in a Collection	6
7.3.1	Using the In-Context Autocomplete Feature in the JSON Editor	9
7.4	Creating Indexes for JSON Collections	10
7.5	Creating Relational Views of JSON Documents	12
7.6	Viewing the JSON Data Guide Diagram for a Collection	13

8 The JavaScript Page

8.1	About the JavaScript Page	1
8.2	Managing MLE Modules	3
8.2.1	Creating an MLE Module	3
8.2.2	About Context Menu Options for MLE Modules	4
8.2.3	Loading Multiple JavaScript Files to Create MLE Modules	5
8.2.4	Example: Creating an MLE Module	6
8.3	Managing MLE Environments	6
8.3.1	Example: Creating an MLE Environment	7
8.4	Managing MLE Call Specifications	8
8.4.1	Example: Creating an MLE Call Specification	10
8.5	Executing JavaScript Code Snippets and Debugging MLE Modules	11
8.5.1	Example: Debugging a JavaScript Module	12

9 The Scheduling Pages

9.1	Overview	2
9.2	Jobs	2
9.2.1	Summary	2
9.2.2	Running	4
9.2.3	History	4
9.2.4	Forecast	4

9.2.5	History (Gantt Chart)	5
9.2.6	Notifications	5
9.2.7	Create or Edit Job	6
9.3	Chains	9
9.3.1	Create or Edit Chain	12
9.3.2	Add or Edit Step	12
9.3.3	Add or Edit Rule	12
9.4	Programs	13
9.4.1	Create or Edit Program	13
9.5	Schedules	14
9.5.1	Create or Edit Schedule	15
9.6	Objects	15
9.6.1	Create Job Class	17
9.6.2	Create File Watcher	17
9.6.3	Create Window	18
9.6.4	Create Window Group	19

Part II Vector Database

Part III Administration

10 The Database Users Page

10.1	Creating or Editing a User	2
10.2	Enabling REST	3
10.3	Deleting a User	3
10.4	Creating a Self Service Schema	4

11 The APEX Workspaces Page

11.1	Create a Workspace	1
------	--------------------	---

12 The Data Pump Page

12.1	Importing Data	2
------	----------------	---

Part IV Monitoring

13 The Monitoring Pages

13.1	The Monitoring Overview Page	1
13.2	The Performance Hub Page	2
13.3	The Instance Viewer Page	3
13.4	The Logins Page	4
13.5	The Alerts Page	5
13.6	The Sessions Page	6
13.7	The Storage Page	7
13.8	The Parameters Page	7
13.9	The Real Time SQL Monitoring Page	8
13.10	The Top SQL Page	10
13.11	The Waits Page	11
13.12	The AWR Page	12

A Supported SQL*Plus and SQLcl Commands in SQL Worksheet

A.1	Supported SQL*Plus Commands	A-1
A.2	Supported SQLcl Commands	A-2
A.2.1	SODA Commands	A-3

Preface

This document provides information about Oracle Database Actions, a web-based interface that provides development and administration features for Oracle Database.

Audience

This online help is intended for those using Oracle Database Actions included with Oracle REST Data Services.

Product Accessibility

Oracle Database Actions provides features to support accessibility. See *Oracle Database Actions Accessibility Guide*.

Related Documents

Oracle REST Data Services Installation and Configuration Guide

Oracle REST Data Services Developer's Guide

[Oracle Autonomous Database](#)

Oracle SQL Developer User's Guide

To download release notes, installation documentation, white papers, or other collateral for SQL Developer, go to the Oracle Technology Network (OTN) at

<http://www.oracle.com/technetwork/>

For the PL/SQL page on OTN, see <http://www.oracle.com/technetwork/database/features/plsql/>

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.
<i>italic</i>	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.

Third-Party License Information

See Third-Party License Information in *Oracle REST Data Services Developer's Guide*

1

Changes in Release 26.2 for *Using Oracle Database Actions*

The new feature is:

- A Vector Database page for managing embedding models, vector tables, and semantic search operations.

See [The Vector Database Page](#)

2

About Oracle Database Actions

Oracle Database Actions is a web-based interface for on-premises Oracle Database that uses Oracle REST Data Services to provide many of the database development and administration features of desktop-based Oracle SQL Developer. The main features include running SQL statements and scripts in the worksheet, exporting data, creating RESTful web services, managing JSON collections, monitoring databases, and creating Data Modeler diagrams.

Note

Some features of Database Actions are only available if you sign in as a user with the DBA role. For such features, a "restricted availability" statement appears at the start of the feature description. For example:

Available only if you signed in as a database user with the DBA and PDB_DBA roles.

Database Actions is also available with:

- Oracle Autonomous Database cloud services. See [Database Actions](#)
- Oracle APEX. See [Accessing SQL Developer Web Directly from APEX](#)

See Also

- [Creating or Editing a User](#) for creating users and assigning roles.
- ORDS System Requirements for system requirements.

2.1 About the Database Actions User Interface

This section describes the Database Actions user interface.

The Database Actions user interface has three components:

- The [Header](#) at the top
- The page body, whose content varies depending on which page you are viewing
- The [Status Bar](#) at the bottom

Header

The header contains the Selector icon, a Search field, the help icon, and the user drop-down list.



- **Selector Icon**

Click the Selector icon  to see the main navigation menu slide into view. Click **Oracle Database Actions** in the header to go to the Launchpad page.

- **Search field**

To enter a search term, click in the Search field or use the shortcut key **Ctrl+K** (**Command+K** for Apple computers). For more information, see [Using the Omnisearch Bar](#).

- **Help Icon**

Click the help icon to open the contextual or online help for the page you are viewing.

- **User Drop-Down List**

The user drop-down list shows the database user you are signed in as, and provides the following items when you open it:

- **Preferences**

The options are:

General

- * **Language:** Select one of the following languages for the user interface: English, German, Spanish, French, Italian, Japanese, Korean, Portuguese, and Chinese
- * **Timezone:** Select **UTC** or **Local time zone** from the drop-down list.
- * **Theme:** Options are Dark, Light or Same as browser. Select **Dark** for changing the display to a dark background (Dark mode).

Code Editor

Provides global settings that apply to the code editor, such as theme, indentation, font family and line numbers.

Worksheet

- * **SQL History:** Controls whether the history of commands executed in the code editor is enabled in the browser or not.


Scripts

Provides options to set the initial state of commands.

- **Log:** Opens a dialog that shows the list of HTTP calls made during your session.
- **About:** Opens a dialog providing version information for the database and other components as well as copyright and licensing information.
- **Sign Out:** Signs you out of your database session.

Status Bar

The status bar contains icons that link to log files. The three icons (Errors, Warnings, Processes) are filters that have been applied to the log file.



Errors, Warnings: Displays an Errors or Warnings dialog, which lists log entries from unsuccessful REST calls or from any other problem in the application.

Processes: Displays a Processes dialog, which logs REST calls that are either finished or ongoing.


Log notification link: Displays a Log dialog, containing log entries of the following types: Errors, Warnings, Processes, SQL History and SQL Result.

2.1.1 Using the Omnisearch Bar

You can access the Omnisearch bar (Search field) in the header from any page in Database Actions.

To enter a search term, click in the search field located at the top right of the header, or use the shortcut key **Ctrl+K** (**Command+K** for Mac OS systems).

In the Omnisearch bar, you can filter the search entry by selecting a category (such as tables, views, templates, and so on) from the displayed list.



If you do select a category, you will then need to enter the exact search term.

If you do not select a category, **ALL** is selected by default. After you enter the search term, a search is performed across all categories. In some cases, the search term is searched across multiple attributes of a category. The following table lists the attributes that are searched for each category:

Category	Attribute
Tables	Name
Views	Name


Category	Attribute
Indexes	Name
Packages	Name
Functions	Name
Procedures	Name
Triggers	Name
Types	Name
Sequences	Name
Charts	Name, Comments, URI_Prefix
Dashboards	Name, Comments, URI_Prefix
Modules	Name, Comments, URI_Prefix
Templates	Comments, URI_Prefix
Handlers	Comments
Roles	Name
Privileges	Label, Name, Comments, Description
OAuth Clients	Name, Description
Database Users	Username, Alias (alias is used depending on user permissions)
APEX Workspaces	Workspace, Workspace display name

Displaying Results for the Search Term

In the results displayed, categories are sorted based on the following criteria:

1. Number of exact matches
2. Most number of matched items

The results within a category are displayed as two groups. The first group consists of items that exactly match the search term. The second group consists of items that contain the search term but do not start with it. Within each group, the items are sorted alphabetically.



You can quickly access previous search terms using the **Recent Searches** list.

2.2 Accessing Database Actions

Oracle Database Actions is included with Oracle REST Data Services.

To access Database Actions:

1. Enable **Database Actions** in Oracle REST Data Services. For more information, see Interactive Command-Line Interface Installation in *Oracle REST Data Services Installation and Configuration Guide*.
2. To use Database Actions, you must sign in as a database user whose schema has been REST-enabled.

Execute the following code as a database user with the DBA role:

```
BEGIN
  ords_admin.enable_schema(
    p_enabled => TRUE,
    p_schema => 'schema-name',
    p_url_mapping_type => 'BASE_PATH',
    p_url_mapping_pattern => 'schema-alias',
    p_auto_rest_auth => TRUE
  );
```

```
commit;  
END;
```

where:

- *schema-name* is the database schema name in all-uppercase.
- *schema-alias* is an alias for the schema name that will appear in the URL the user will use to access Database Actions. Oracle recommends that you do not use the schema name itself as a security measure to keep the schema name from being exposed.
- *p_auto_rest_auth* specifies that the REST /metadata-catalog/ endpoint requires authorization. REST uses the metadata-catalog to get a list of published services on the schema.

3. Sign in to Database Actions. See [Signing In to Database Actions](#)

2.3 Signing In to Database Actions

You can sign in to Database Actions by connecting to a single default database or to multiple databases or by using a schema alias.

Note

- From Oracle REST Data Services, you can access the Database Actions Sign In page from the ORDS landing page (<https://example.com/ords>).
- From APEX, you can open Database Actions directly from the SQL workshop dropdown menu. See [Accessing SQL Developer Web Directly from APEX](#)

The Sign in page has a languages menu on the top right where you can select the language to display.

Default Database

To connect to a single default database:

- In the Database Actions Sign-in page, enter your user name and password and click **Sign in**.

Note

In the Password field, if the CAPS LOCK key is turned on, an up arrow icon is displayed in the right corner of the field.

ORDS Schema Alias

To connect to a database using an ORDS schema alias, which is different from your user name:

- In the Database Actions Sign-in page, expand **Advanced** and in the Path field, enter the schema alias.

See [Enabling REST](#) to know how to set up a schema alias.

- Enter your user name and password and click **Sign in**.

Multiple Databases

To connect to a database in a database pool:

- In the Database Actions Sign-in page, expand **Advanced** and in the Path field, enter the database pool name specified while installing Oracle REST Data Services and the schema name, <database pool name>/<schema name>.

See Configuring Oracle REST Data Services for Multiple Databases in *Oracle REST Data Services Installation and Configuration Guide* for more information about database pools.

- Enter your user name and password and click **Sign in**.

2.4 The Home Page - Launchpad

The Launchpad page for Database Actions can now dynamically display categories and features based on whether you are logged in to an on-premises, Autonomous AI Database or OCI environment. For on-premises, the main categories displayed are Pinned and Recently Visited, Development, Administration and Monitoring. For Autonomous AI Database, additional categories displayed are Data Studio and Downloads. Descriptions for features belonging to these categories are provided below.

The first time you log in, the Development tab is displayed by default. All features that pertain to the Development category are displayed vertically on the left side of the screen. When you hover over a feature name, its description is displayed on the right along with links to specific documentation pages. A new feature is highlighted with a blue label.

In Pinned & Recently Visited, Pinned lists all features that you have pinned (saved). To pin a feature, click the Pin icon displayed on the top right of the feature description. There is no limit on the number of features that can be pinned. **Recently Visited** lists the last seven features that you have accessed.

To navigate to a specific feature such as the REST pages, click **REST** on the left or click **Open** in the feature description part.

The following is a description of each feature displayed on the home page.

Development

- **SQL**: Enter and execute SQL and PL/SQL commands, and create database objects. See [The SQL Page](#)
- **Data Modeler**: Create diagrams from existing database schemas, generate DDL statements, and create reports. See [The Data Modeler Page](#)
- **APEX**: Link to the Oracle Application Express sign-in page. Application Express is a rapid web application development platform for the Oracle database. See Oracle Application Express documentation
- **REST**: Develop RESTful web services and ensure secure access. See [The REST Pages](#)
- **Liquibase**: View changelogs for your schema. See [The Liquibase Page](#)
- **Charts**: Create charts and dashboards containing multiple charts using SQL queries. See [The Charts and Dashboards Page](#)
- **JSON**: Manage and query JSON collections. JSON is available only if you are signed in as a database user with the SODA_APP role. See [The JSON Page](#)
- **JavaScript**: Execute JavaScript code in Oracle Database. See [The JavaScript Page](#)

- **Scheduling:** Provide details of scheduled jobs, chains, programs and schedules. See [The Scheduling Pages](#)

Vector Database

Provides a centralized way to manage vector-specific operations. See [The Vector Database Page](#)

Administration

Administration is available only if you are signed in as a database user with administrator rights.

- **Database Users:** Perform user management tasks such as create, edit, and REST enable users. Users have access to features based on user privileges. See [The Database Users Page](#)
- **Data Pump:** Monitor Data Pump jobs initiated through the available Database API endpoints, the DBMS_DATAPUMP package, or the SQL Developer Data Pump Export and Import wizards. See [The Data Pump Page](#)
- **APEX Workspaces:** Create and manage APEX workspaces.

Monitoring

Monitoring is available only if you are signed in as a database user with administrator rights.

- Monitor database activity and performance using various tools.

Monitoring is available to database users with DBA and PDB_DBA roles. The only exception is the Real Time SQL Monitoring feature, which is also available to non-administrator users for Oracle Database 19c and later versions.

See [The Monitoring Pages](#)

Part I

Development

This part provides information about the following topics:

Topics:

- [The SQL and Data Modeler Pages](#)
- [The REST Pages](#)
- [The Liquibase Page](#)
- [The Charts and Dashboards Page](#)
- [The JSON Page](#)
- [The Scheduling Pages](#)

3

The SQL and Data Modeler Pages

Use the SQL page to enter and execute SQL and PL/SQL statements and create database objects. The Data Modeler page enables you to create diagrams from existing schemas, retrieve data dictionary information, generate DDL statements, and export diagrams.

Topics

- [The Overview Page](#)
- [The SQL Page](#)
- [The Data Modeler Page](#)
- [Creating and Editing Database Objects](#)

3.1 The Overview Page

The Overview page contains widgets that provide a general overview of the activity in the SQL and Data Modeler pages.

To navigate to the Overview page, click **Selector**  and then under Development, select **Overview**.

- **My Worksheets:** Displays your saved worksheets. You can click the name of the worksheet to open it in the Worksheet page.
- **My Diagrams:** Displays the Data Modeler diagrams that have been saved. You can click the name of the diagram to open it in the Data Modeler page.
- **Recently Modified Objects:** Displays a timeline of the created, modified, and dropped objects in the database. You can zoom in and out using the + and – icons. You can also move horizontally by dragging the cursor to the right or left.
- **Invalid Objects:** Displays the invalid objects in your schema.
- **Table Stats Freshness:** Displays the time period since the tables were last analyzed.

You can right-click the header in Invalid Objects, Table Stats Freshness, My Worksheets, or My Diagrams to manage or sort columns:

- **Columns:** Enables you to select columns to show or hide.
- **Sort:** Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

Right-click the body of the display table to count rows or to view records:

- **Count Rows:** Displays the number of rows in the table.
- **Single Record View:** Enables you to view data for a table or view, one record at a time.
- **Copy:** Copies data from a cell or a row or a range of rows. To copy from more than one row, select the rows you want to copy, right-click by pressing the SHIFT or CTRL key, and select **Copy**.

3.2 The SQL Page

The SQL page enables you to enter and execute SQL and PL/SQL statements, and create database objects.

To navigate to the SQL page, do either of the following:

- In the Launchpad page, select the **Development** tab and click **SQL**.
- Click **Selector**  to display the navigation menu. Under Development, select **SQL**.

You can use SQL and PL/SQL statements in the worksheet to create a table, insert data, create and edit a trigger, select data from a table, and save that data to a file. Some other features are syntax highlighting and error detection.

The SQL page consists of the left pane for navigating worksheets and objects, the editor for executing SQL statements, and the output pane for viewing the results. These panes are described in the following sections:


- [Object Navigator and Files](#)
- [Executing SQL Statements in the Code Editor](#)
- [Viewing the SQL Output](#)
- [Loading Data](#)

3.2.1 Object Navigator and Files

The Navigator tab in the left pane displays saved objects for the selected schema. The Files tab enables you to view and open files saved in your browser or local system.

The following figure shows the various elements in the left pane.

Figure 3-1 Left Pane in SQL



Navigator Tab

Displays saved objects for the selected schema.

- **Schema and Object Type selector:** Use the drop-down lists to select the schema and filter the results by object type.
- **Search:** Searches the contents of a saved worksheet or search for objects in the Navigator tab by name. The search functionality is not case-sensitive, retrieves all matching entries, and does not require the use of wild card characters.
- **Context menu:** Options in the context menu are:

- **Open** to browse properties and data relevant to the object type (tables and views).

The Data pane displays the data for a table, view, or materialized view.

To edit an entry, double-click a cell to make edits. You can also click  and enter the value. When you make an edit, the border of the gutter cell in that row changes to blue.

The icons available in the Data pane are:


- * **Create Row:** Insert a new row into the database table. When you insert a row using the **Create Row** icon, the row is committed into the database.
- * **Delete Selected:** Mark the selected row for deletion. When you mark a row for deletion, the border of the row changes to red.
- * **Commit:** Commit all changes made to the database.

- * **Undo All:** Revert all changes that are marked for commit.

Use the **Filter** icon at the top right corner to filter the column data. You can also right-click a cell to access the context menu to count rows, view a single record, export or copy the cell text to clipboard.

To view a Binary Large Object data type (BLOB), click the pencil icon for a BLOB data type. In the View Value dialog:

- * The **Image** tab displays the loaded image, if the loaded BLOB type is an image.
- * The **Text** tab displays the text file, if the loaded BLOB type is text .
- * The **Information** tab displays the details and allows you to perform the following actions:
 - * **Download:** To download the image/text file of BLOB data type .
 - * **Load:** To insert an image/text of BLOB data type.
 - * **Set NULL:** To delete the object and set the value as NULL.



- **Edit** edits the properties of an existing object.
- **Add** creates an object based on the object type selected.
- **Use as Template** creates an object by using the properties of an existing object as the template.
- **REST**
 - * **Enable** enables REST access for the database object. See [Enabling REST Access for a Database Object](#)
 - A REST enabled object is indicated by a **REST Enabled** icon  in the Navigator tab.
 - * **Disable** disables REST access for the database object after it is enabled. See [Disabling REST Access for a Database Object](#)
 - * **cURL Command** generates a cURL call for a selected HTTP method for the database object. See [Generating cURL Requests for a REST-Enabled Database Object](#)
- **Data Loading** loads data from local files into a table.
- **Quick DDL** generates Data Definition Language statements for the object.


- **Refresh** : Refreshes the objects or worksheets listed in the left pane.
- **Object Submenu** : Opens the Create Object dialog to create a new object based on the type selected in the drop-down list.
- **Help** : Provides contextual help documentation.

Files Tab

Enables you to open files from your browser or local device.

Note

The DEVICE category in the left pane is displayed only when using a Chromium-based browser in a secure context (HTTPS).




All files: Use the drop-down list to filter files by browser or device.

The context menu options for a file are **Open** and **Delete**. In Device, the corresponding option for deleting a file is **Forget**. In this case, the file is not deleted, instead the reference to the file is removed.

Drag and Drop Objects and Files into the Worksheet

You can drag objects from the left pane and drop them into the worksheet editor in the right pane.

- If you drag and drop a table or view, you are prompted to select one of the following SQL statements: Insert, Update, Select, or Delete. For example, if you choose Select, a Select statement is constructed with all columns in the table or view. You can then edit the statement, for example, modifying the column list or adding a WHERE clause.

Figure 3-2 Insert Select Query

If you choose Object Name, the name of the object prefixed by the schema name is added to the worksheet.

- If you drag and drop a function or procedure, you can choose to insert the name or the PL/SQL code of the function or procedure in the worksheet. If you select the PL/SQL code, you can enter the parameters before inserting the code into the worksheet.

3.2.2 Executing SQL Statements in the Code Editor

The code editor in the SQL page enables you to run SQL statements, PL/SQL scripts, and JavaScript code. The main features include in-context code completion, syntax highlighting, and error debugging.

You can enter SQL and PL/SQL statements to specify actions such as creating a table, inserting data, selecting data, or deleting data from a table. SQL keywords are automatically highlighted. For multiple statements, you must terminate:


- Each non-PL/SQL statement with either a semicolon (;) or a slash (/) on a new line
- Each PL/SQL statement with a slash (/) on a new line

The **PL/SQL editor** is triggered when opening the following object types: Functions, Procedures, Packages and Types. This editor helps you detect errors in your PL/SQL code during compilation. The output includes details such as the specific line and column where the error is detected, along with a link to go to the relevant position in the code block.

The **JavaScript worksheet mode** supports the Multilingual Engine syntax in Oracle Database release 21c. For more details, see [Support for Multilingual Engine](#).


For SQL*Plus and SQLcl statements supported in the worksheet, see [Supported SQL*Plus and SQLcl Commands in SQL Worksheet](#).

If you press **Ctrl+Space**, the editor provides you with a list of possible completions at the insertion point that you can use to autocomplete code that you are editing. This list is based on the code context at the insertion point.



The editor offers a comprehensive list of commands available through the Command Palette. To open the Command Palette, press **Ctrl+Shift+P**. For a list of keyboard shortcut keys, see [Keyboard Shortcuts](#).

An error in the code is signified by a red dot in the left gutter and a squiggle line beneath the specific text. When you hover over it, you see a pop-up displaying possible fixes for resolving the error.



You can set code editor preferences using the **Preferences** option available in the top-right user drop-down list. Some of the code editor options that you can customize are Theme (Light, Dark and High contrast dark), Font size and family, Tab size, Word wrap, Ruler, Line numbers and so on.

The SQL toolbar contains icons for the following operations:

- **Worksheet** drop-down list
 - **Open:** Opens a file from the browser or device.
 - **Open Recent:** Displays the recently accessed files. If there are more than five files in the recent list, then a **More** link is displayed for viewing the additional files.
 - **Save As:** Saves a file to the browser or device.
- **New File**

- Select **Worksheet** to create a worksheet.
- Select **PL/SQL** to create a PL/SQL object type. The editor switches to a PL/SQL mode.
- Select **JavaScript** to create a JavaScript file. The editor switches to a JavaScript mode and **(JS)** is added to the file name.
- **Open** opens a file from your browser or device. To open a file from your device, in the Open File slider, click **Open File** and browse to select the file, or drag and drop the file into the slider.
- **Run Statement**
 - **Run Statement** executes the selected statements or the statement at the mouse pointer in the worksheet editor. The SQL statements can include bind variables and substitution variables of type VARCHAR2 (although in most cases, VARCHAR2 is automatically converted internally to NUMBER if necessary). A dialog box is displayed for entering variable values.
 - **Run Statement Without Pagination** executes the selected statements without wrapping it inside a ROW_NUMBER() OVER analytic function to implement paging. This means the query is run 'as is' in the worksheet, and only fetches the first 256 rows.
- **Run Script**
 - **Run as SQL Script** executes all statements in the worksheet editor using the Script Runner. The SQL statements can include bind variables (but not substitution variables) of type VARCHAR2 (although in most cases, VARCHAR2 is automatically converted internally to NUMBER if necessary). A dialog box is displayed for entering bind variable values.
 - **Run as JavaScript** executes the code as a JavaScript file. This option is used only with JavaScript code. If the JavaScript code is added to a PL/SQL block, then select "Run as SQL Script" to execute the script.
- **Compile** (for PL/SQL toolbar) performs a PL/SQL compilation of the subprogram.
- **Create Chart** creates a chart for the corresponding SQL statement entered in the editor. In contrast, you can partially highlight a subquery and create a chart. A slider window is displayed for entering the chart parameters. For a description of the fields, see [Creating or Editing a Chart](#).

If the SQL statement is syntactically incorrect or incomplete, an error/warning notification is displayed.
- **Explain Plan** generates the execution plan for the statement (internally executing the EXPLAIN PLAN statement). The execution plan is automatically displayed in the Explain Plan tab in the worksheet output pane. See [Viewing the SQL Output](#)
- **Autotrace** runs the statement and collects run time statistics and the actual execution plan. The Autotrace output is displayed in the Autotrace tab in the worksheet output pane. Currently, there are no preferences available.
- **Download Editor Content** downloads the content of the worksheet as a SQL file to the local system.
- **Format** formats the SQL statement in the editor, such as capitalizing the names of statements, clauses, keywords, and adding line breaks and indentation.
- **Clear** removes the statements from the editor.
- **Tour** provides a guided tour of the worksheet highlighting salient features and providing information that is useful if you are new to the interface.
- **Help** provides context-related help and provides a link to the help documentation.

- **Open in Fullscreen** opens the editor in full screen mode.

3.2.2.1 About Session State in Database Actions

In Database Actions, a connection to the database is stateless.

In a stateless environment, each HTTPS request from a client maps to a new database session. Therefore, a session begins and ends with every SQL statement or script execution.

As the session state is not maintained, session attributes do not persist and commands such as ROLLBACK and COMMIT do not apply. If a SQL statement or script executes successfully, an implicit commit is performed. If it executes with an error, an implicit rollback is performed.

Therefore, when needed, include the ROLLBACK and COMMIT commands or session attributes in the PL/SQL code block that is sent to the database for a session.

The only configuration commands that persist during a session in Database Actions are:

- SET DEF[INE] <ON|OFF|*prefix_character*>
- SET ESC[APE] <ON|OFF|*escape_character*>
- SET TIMI[NG] <ON|OFF>

3.2.2.2 Keyboard Shortcuts

This section lists the keyboard shortcuts for various commands in the SQL page.

Table 3-1 Keyboard Shortcuts

Windows	MacOS	Description
Ctrl + Enter	Cmd + Enter	Runs the code as query.
Alt + Tab	Option + Tab	Focus next element.
Ctrl+ Esc / Escape	Cmd + Esc / Escape	Remove focus from editor.
Ctrl + Down Arrow	Cmd + Down Arrow	Moves to the next SQL code from history.
Ctrl + Up Arrow	Cmd + Up Arrow	Moves to the previous SQL code from history.
Ctrl + S	Cmd + S	Saves the current worksheet.
Ctrl + O	Cmd + O	Opens the worksheet browser dialog.
Ctrl + I	Cmd + I	Downloads the content of the editor.
F1	Fn + F1	Opens the help topic.
Shift + Esc	Shift + Esc	Focus previous element.
F5	Fn + F5	Runs code as script.
F6	Fn + F6	Shows Autotrace.
F10	Fn + F10	Shows Explain Plan.
F11	Fn + F11	Creates a chart.
Ctrl + B	Cmd + B	Opens the "Convert Case" drop-down list.
Ctrl + F7	Cmd + Fn + F7	Formats code in the editor.
Ctrl + Space	Ctrl + Space	Autocompletes code (shows hints).

The following table lists the keyboard shortcuts for commands in the Command Palette.

Table 3-2 Command Palette Keyboard Shortcuts

Windows	MacOS	Description
Ctrl + Alt + Up	Cmd + Option + Up	Add Cursor Above
Ctrl + Alt + Down	Cmd + Option + Down	Add Cursor Below
Shift + Alt + I	Shift + Option + I	Add Cursors to Line Ends
Ctrl + K Ctrl + C	Cmd + K Cmd + C	Add Line Comment
Ctrl + D	Cmd + D	Add Selection To Next Find Match
Shift + Alt + Down	Shift + Option + Down	Copy Line Down
Shift + Alt + Up	Shift + Option + Up	Copy Line Up
Ctrl + U	Cmd + U	Cursor Undo
-	Cmd + Backspace	Delete All Left
-	Ctrl + K	Delete All Right
Ctrl + Shift + K	Shift + Cmd + K	Delete Line
Shift + Alt + Right	Shift + Ctrl + Cmd + Right	Expand Selection
Ctrl + F	Cmd + F	Find
Enter	Enter	Find Next
Ctrl + F3	Cmd + Fn + F3	Find Next Selection
Shift + Enter	Shift + Enter	Find Previous
Ctrl + Shift + F3	Shift + Cmd + Fn + F3	Find Previous Selection
-	Cmd + E	Find With Selection
Ctrl + Shift + [	Option + Cmd + [	Fold
Ctrl + K Ctrl + 0	Cmd + K Cmd + 0	Fold All
Ctrl + K Ctrl + /	Cmd + K Cmd + /	Fold All Block Comments
Ctrl + K Ctrl + 8	Cmd + K Cmd + 8	Fold All Regions
Ctrl + K Ctrl + 1	Cmd + K Cmd + 1	Fold Level 1
Ctrl + K Ctrl + 2	Cmd + K Cmd + 2	Fold Level 2
Ctrl + K Ctrl + 3	Cmd + K Cmd + 3	Fold Level 3
Ctrl + K Ctrl + 4	Cmd + K Cmd + 4	Fold Level 4
Ctrl + K Ctrl + 5	Cmd + K Cmd + 5	Fold Level 5
Ctrl + K Ctrl + 6	Cmd + K Cmd + 6	Fold Level 6
Ctrl + K Ctrl + 7	Cmd + K Cmd + 7	Fold Level 7
Ctrl + K Ctrl + [	Cmd + K Cmd + [	Fold Recursively
Ctrl + Shift + \	Shift + Cmd + \	Go to Bracket
Ctrl + G	Ctrl + G	Go to Line...
Alt + F8	Option + Fn + F8	Go to Next Problem(Error, Warning, Info)
F8	Fn + F8	Go to Next Problem in Files (Error, Warning, Info)
F7	Fn + F7	Go to Next Symbol Highlight
Shift + Alt + F8	Shift + Option + Fn + F8	Go to Previous Problem (Error, Warning, Info)
Shift + F8	Shift + Fn + F8	Go to Previous Problem in Files (Error, Warning, Info)
Shift + F7	Shift + Fn + F7	Go to Previous Symbol Highlight
Ctrl +]	Cmd +]	Indent Line
Ctrl + Shift + Enter	Shift + Cmd + Enter	Insert Line Above

Table 3-2 (Cont.) Command Palette Keyboard Shortcuts

Windows	MacOS	Description
-	Ctrl + J	Join Lines
Ctrl + K Ctrl + D	Cmd + K Cmd + D	Move Last Selection To Next Find Match
Alt + Down	Option + Down	Move Line Down
Alt + Up	Option + Up	Move Line up
F1 (All browsers)	Fn + F1	Open Command palette
Ctrl + Shift + P (Google Chrome only)		
Ctrl + [	Cmd + [	Outdent Line
Ctrl + K Ctrl + U	Cmd + K Cmd + U	Remove Line Comment
Ctrl + H	Option + Cmd + F	Replace
Ctrl + Shift + .	Shift + Cmd + .	Replace with Next Value
Ctrl + Shift + ,	Shift + Cmd + ,	Replace with Previous Value
Ctrl + Shift + L	Shift + Cmd + L	Select All Occurrences of Find Match
Alt + F1	Option + Fn + F1	Show Accessibility Help
Shift + F10	Shift + Fn + F10	Show Editor Context Menu
Ctrl + K Ctrl + I	Cmd + K Cmd + I	Show Hover
Shift + Alt + Left	Shift + Ctrl + Cmd + Left	Shrink Selection
Shift + Alt + A	Shift + Option + A	Toggle Block Comment
Ctrl + K Ctrl + L	Cmd + K Cmd + L	Toggle Fold
Ctrl + /	Cmd + /	Toggle Line Comment
Ctrl + M	Shift + Ctrl + M	Toggle Tab Key Moves Focus
-	Shift + Ctrl + T	Transpose Letters
Ctrl + Space	Ctrl + Space	Trigger Suggest
Ctrl + K Ctrl + X	Cmd + K Cmd + X	Trim Trailing Whitespace
Ctrl + Shift +]	Option + Cmd +]	Unfold
Ctrl + K Ctrl + J	Cmd + K Cmd + J	Unfold All
Ctrl + K Ctrl + 9	Cmd + K Cmd + 9	Unfold All regions
Ctrl + K Ctrl +]	Cmd + K Cmd +]	Unfold Recursively

3.2.2.3 Support for Multilingual Engine

Database Actions provides support for Multilingual Engine (MLE) by enabling you to run JavaScript code in the worksheet.

Prerequisites

For the availability of MLE features in Database Actions, you need the:

- DBMS_MLE package in Oracle Database Release 21c and later versions. For more information, see DBMS_MLE in *Oracle Database PL/SQL Packages and Types Reference*.
- EXECUTE_DYNAMIC_MLE and EXECUTE_ON_JAVASCRIPT privileges assigned to you.

You can work with JavaScript code in the worksheet in the following ways:

- Create a JavaScript worksheet

- Execute JavaScript code in a standard worksheet
- Execute JavaScript code as a PL/SQL block


Create a JavaScript Worksheet

You can open the worksheet in JavaScript mode. You will also see the toolbar icons and output tabs change to reflect the JavaScript mode.

To create and save a JavaScript worksheet:

1. In the toolbar, expand the **New File** icon and select **JavaScript**.
2. When you enter code in the worksheet, the JavaScript code is automatically highlighted.
3. Execute the JavaScript code in the worksheet using **Run Script**.
4. Click **Save**.

In the Files pane, **(JS)** is added to the JavaScript file name. This enables you to quickly identify the JavaScript file.




Execute JavaScript code in a standard worksheet

In a standard worksheet (when the worksheet is not in JavaScript mode):

1. Select the JavaScript code to execute.
2. In the worksheet toolbar, expand the **Run Script** icon and select **Run as JavaScript**.

Note


Selecting **Run as SQL script** results in an error.



Execute JavaScript code as a PL/SQL block

You can use a PL/SQL code block to execute JavaScript code.

From the worksheet toolbar, expand the **Run Script** icon and select **Run as SQL script (F5)**.




3.2.3 Viewing the SQL Output

The lower right pane in SQL displays the output of the operation executed in the SQL editor.

The following figure shows the output pane in the SQL page.

Figure 3-3 Output Pane



The screenshot shows the Oracle SQL Developer interface. The top toolbar includes icons for file operations, execution, and formatting. The main workspace displays a SQL query in a worksheet:

```

1 SELECT
2   DEPARTMENT_ID,
3   DEPARTMENT_NAME,
4   MANAGER_ID,
5   LOCATION_ID
6 FROM
7   HR.DEPARTMENTS;
8

```

Below the worksheet, the Output Pane is visible with the following tabs: Query Result, Script Output, DBMS Output, Explain Plan, Autotrace, SQL History, and Data Loading. The 'Query Result' tab is selected, showing the results of the query in a table. The table has five columns: an index, department_id, department_name, manager_id, and location_id. The results are as follows:

	department_id	department_name	manager_id	location_id
1	10	Administration	200	1700
2	20	Marketing	201	1800
3	30	Purchasing	114	1700
4	40	Human Resources	203	2400

The output pane has the following tabs:

- **Query Result:** Displays the results of the most recent Run Statement operation in a display table.
- **Script Output:** Displays the text output from your statements executed as a script using the script engine.
- **DBMS Output:** Displays the output of DBMS_OUTPUT package statements.
- **Explain Plan:** Displays the plan for your query using the Explain Plan command. The default view is the diagram view. For more information, see [Using the Explain Plan Diagram](#).
- **Autotrace:** Displays the session statistics and execution plan from v\$sql_plan when executing a SQL statement using the Autotrace feature. Displays the output if you clicked the Autotrace icon.
- **SQL History:** Displays the SQL statements and scripts that you have executed. To re-enter a previously executed query in the worksheet, double-click the query in the history list. You can search for specific statements by clicking the Search icon. The Search functionality is case-sensitive, retrieves all entries that contain the search text, and does not require wildcard characters.

Note

The executed statements are saved in the database, and are therefore available across all devices. For a read-only database, the SQL statements are saved on your browser.

- **Data Loading:** Displays a report of the total rows loaded and failed for all visible tables (including tables from other schemas).

The icons in this pane are:

- **Clear output:** Clears the output.
- **Show info:** Displays the SQL statement for which the output is displayed.
- **Open in new tab:** Opens the query result or explain plan in a new window.
- **Share Query Results:** This is only applicable to the Query Result and is available for use only when a query has been successfully executed, generating result data. If no result data is available, then this option is disabled for use.

It enables you to share query results with others using a shareable URL to the query results. On clicking the Share button, a URL is generated. By simply copying it to your clipboard, you can easily share the results with non-authenticated users. Therefore, anyone you provide this URL to can access it without authentication (like a user name/password), and for the time limit you define using **Expiration date**. If you don't specify expiration date and time, then it will have a default expiration date as the current date and a default minimum expiration time of one hour. Although a default minimum expiration duration is set, there is no maximum limit on how long the expiration duration can be set. This means that you can set the expiration duration as long as you like as per your specific business requirements.

Note

Query Results that you share using this option are based on a specific query that is executed at a particular point in time. If the underlying objects or views referenced in the query are modified later, the query result data may change or vary.

- **Download:** This is applicable only for Query Result. Enables you to download the query result to your local computer in CSV, JSON, XML, or TEXT (.tsv) format.

In the Query Result tab, in the display table, the context menu (right-click) for the row header consists of the following:

- **Columns** enables you to select columns to hide.
- **Sort** displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

Figure 3-4 Context Menu for Row Header

OWNER	OBJECT_NAME	SUBOBJECT_NAME
PUBLIC	STM	(null)
OUTLN	OL\$	(null)
OUTLN	OL\$NAME	(null)

The context menu for the rest of the display table consists of the following commands:

- **Count Rows** displays the number of rows in the result set for your query.
- **Single Record View** enables you to view data for a table or view, one record at a time.


- **Export** generates the file for download based on the format selected, which can be XML, CSV (comma-separated values including a header row for column identifiers), Insert , Delimited, Fixed, HTML, JSON, or TEXT.
 - **Format:** Select the format to export from the drop-down list.
 - **Line Terminator:** Identifies the terminator for each line. The line terminator is not included in the data exported. If the preview page shows the data in one single row, the correct terminator is not specified.
 - **Header:** Controls whether the first row is a header row or the first row of data.
 - **Left and Right Enclosure:** Enclosures are used for character data and are optional. Enclosures are not included in the data exported.

Note

If a popup blocker is enabled, it will prevent the file from downloading.

- **Copy** copies data from a cell or a row or a range of rows.

Figure 3-5 Context Menu



3.2.3.1 Using the Explain Plan Diagram

The Explain Plain diagram view is a graphical representation of the contents of PLAN_TABLE, which is the default table for results of the EXPLAIN PLAN statement. The hierarchical nature of the steps in the execution plan is depicted in the diagram.


By default, three levels of steps are visible in the diagram. You can use the +/- signs at the bottom of each step (available when the step has children) to expand or collapse. To view all

steps in the diagram, use  **Expand All** in the toolbar.

The diagram also provides the following details:


- Cardinality (number on the arrow to the parent step), which is the number of rows processed
- Operation and options applied in that step
- Execution order, which is the sequential number in the order of execution
- Access predicates CPU cost in percentage (orange bar)
- Total CPU cost for the step in percentage (blue circle)
- Estimated I/O Cost, Bytes processed and Cost metrics

You can see a brief description pop-up when you hover over any of these statistics in a step.



The icons in the toolbar are:

- **Advanced View/Diagram View:** Switches between Advanced and Diagram views. The Advanced View displays data from PLAN_TABLE in a mixed tabular/tree view. The Diagram View displays the data in a flowchart format.
- **Print Diagram:** Prints the diagram.
- **Save to SVG:** Saves the diagram to file in SVG format.
- **Zoom In, Zoom Out:** If a step is selected in the diagram, clicking the Zoom In icon ensures that it remains at the center of the screen.
- **Fit Screen:** Fits the entire diagram in the visible area.
- **Actual Size:** Sets the zoom factor to 1.
- **Expand All:** Displays all steps in the diagram.
- **Reset Diagram:** Resets the diagram to the initial status, that is, only three levels of steps are displayed.
- **Show Info:** Shows the SELECT statement used by the Explain Plan functionality.
- **Open in New Tab:** Opens the diagram view in a new tab for better viewing and navigation. The diagram is limited to the initial SELECT statement.
- **Min Visible Total CPU Cost(%):** Defines the threshold to filter steps with total CPU cost less than the the provided value.
Enter a value between 0 and 100. There is no filtering for 0.
- **Plan Notes:** Displays the Explain Plan notes.



Properties

Double-click or press **Enter** on a selected step to open the Properties slider, which provides more information about that step. See `PLAN_TABLE` in *Oracle Database Reference* for a description of each property.

The Properties slider shows:

- All information for that step extracted from `PLAN_TABLE` in a tabular format. Nulls are excluded.
You can select **JSON** to view the properties in JSON format.
- Information from `OTHER_XML` column of `PLAN_TABLE`.
The information is displayed in JSON format.


Navigation

- Press the **Tab** key to move through the steps in the execution order. The selected step has a blue border around it.
To move in the reverse direction, press the **Shift + Tab** keys.
If no step is selected, pressing the **Tab** key selects the step with execution number 1.
- Depending on the zoom level, use horizontal and vertical scrollbars to view different parts of diagram.
Click the left mouse button and hold it to pan the diagram around up and down.



Use the icon at the bottom right to scroll to the top of the diagram.

- The Diagram Navigator is at the top right corner and represents a smaller copy of the diagram on a grey background. The rectangle border allows zoom-in and zoom-out operations and moves to show different parts of the diagram.



The diagram navigator shows a list with steps having more than 1% CPU cost in descending order. Click a step in the list to navigate to the same step in the diagram, enabling you to see it in the context of the other steps.

3.2.4 Loading Data

In the SQL page, you can load data from one or more local files into one or more tables.

The file formats that you can load are CSV, XLS, XLSX, TSV, TXT, XML, JSON, and Newline Delimited JSON (NDJSON).


- [Loading Data from a Local File to a New Table](#)
- [Loading Data from a Local File to an Existing Table](#)
- [Loading Data from Multiple Local Files into Multiple Tables](#)

3.2.4.1 Loading Data from a Local File to a New Table

To load data from a local file to a new table:

1. You can start in one of the following ways:
 - In the Navigator tab, in the left pane, click **Object submenu** **...**, select **Data Loading**, and then select **Upload Data into New Table**.

Figure 3-6 Upload Data to New Table Option



- In the Navigator tab, drag and drop the local file into the left pane. When you drag a file into the pane, the following message is displayed: Drop the file here to start.

The Upload Data into New Table is displayed. A preview of the data is displayed in a grid format.

2. Click **Show/Hide options**  to display options that you can modify for data preview:
- **Column names:** Select **Get from file** to display column headers in the first row.
 - **Encoding:** An option to select the encoding type is visible when the loaded file is in plain text format (CSV, TSV, or TXT). The default encoding type is UTF-8.
 - **Text enclosure** and **Field delimiter:** These options are visible only when the selected file is in plain text format (CSV, TSV, or TXT). Select or enter the character used in the source file for text enclosure and field delimiter.
 - **Rows to skip:** Enter or use the up and down arrows to select the number of rows to skip.
 - **Preview size:** Enter or use the up and down arrows to select the number of rows to preview.
 - **Limit rows to upload:** If you select this option, you need to specify the rows to load. Use the up and down arrows to select the number of rows to load.

To remove the options selected and the data preview, click **Clear**.

After selecting the required options, click **Apply**, and then click **Next**.

3. In Table Definition, do the following:
- In the **Table Name** field, enter a name for the target table.
 - Select the check box at the beginning of a row to add the column to the target table.
 - Select or enter values for column attributes such as Column Name, Column Type, Precision, Scale, Default, Primary Key and Nullable.
 - The Format Mask column appears for date, timestamp and numeric types of data. For date and timestamp types, you must select a value from the drop-down list or type the value in the Format Mask field. For numeric types, the format mask is optional.

For a date and timestamp column, you need to supply a compatible format mask that describes the data being uploaded. For example, if the date data looks like 12-FEB-2021 12.21.30, you need to supply a date mask of DD-MON-YYYY HH.MI.SS. The format mask is automatically determined based on the data in the file. You need to review the suggested format mask and if needed, modify it by entering the format directly into the target cell.

Figure 3-7 Table Definition Step in Upload Data into New Table

Upload Data Into New Table

1 Data preview 2 Table definition 3 Review

File: employees.csv Schema: PDBDBA Table name: EMPLOYEES

☒	Column Name	Column Type	Length/Precision	Scale	Default	Primary Key	Nullable	Format mask
☒	EMPLOYEE_NO	NUMBER				☐	☒	
☒	NAME	VARCHAR2	4000			☐	☒	
☒	EMAIL	VARCHAR2	4000			☐	☒	
☐	PHONE_	NUMBER				☐	☒	
☒	HIRE_DATE	TIMESTAMP WITH TIM...	4000			☐	☒	DD-MM-RRRR
☒	LAST_NAME	VARCHAR2	4000			☐	☒	

Click **Next**.

- Review the generated DDL code based on the selections made in the previous screens. The mapping of the source to target columns are also displayed.

Click **Finish**. After the data is successfully loaded, the new table is displayed in the Navigator tab.

- For a detailed report of the total rows loaded and failed, do one of the following:
 - Right-click the table in the Navigator tab, select **Data Loading**, and then select **History**. This displays the report for a specific table.
 - In the Navigator tab, select Object submenu **...**, select **Data Loading**, and then select **History**. This displays the report for all tables in the schema that is selected in the Navigator tab.
 - In the worksheet output pane, select the **Data Loading** tab. This displays the report for all visible tables (including tables from other schemas).

A summary of the data loaded is displayed in the History dialog. If any data failed to load, you can view the number of failed rows in the Failed Rows column. Click the failed rows column to open a dialog showing the failed rows.


In the History dialog, you can also search for files loaded by schema name, table name, or file name. To remove the loaded files, click Remove all history .

You can also load data from a file to a new table using the steps in [Loading Data from Multiple Local Files into Multiple Tables](#).

3.2.4.2 Loading Data from a Local File to an Existing Table

To load data from a local file to an existing table:

- In the Navigator tab, in the left pane, right-click the table that you want to load data into, select **Data Loading**, and then select **Upload Data**.

Figure 3-8 Upload Data Option for an Existing Table

The Import data dialog is displayed.

2. Drag and drop the file from your system into the dialog, or click **Select Files** to browse for the file and open it.

A preview of the data is displayed in a grid format.

3. Click **Show/Hide options**  to display options that you can modify for data preview:
 - **Column names:** Select **Get from file** to display column headers in the first row.
 - **Encoding:** An option to select the encoding type is visible when the loaded file is in plain text format (CSV, TSV, or TXT). The default encoding type is UTF-8.
 - **Text enclosure** and **Field delimiter:** These options are visible only when the selected file is in plain text format (CSV, TSV, or TXT). Select or enter the character used in the source file for text enclosure and field delimiter.
 - **Rows to skip:** Enter or use the up and down arrows to select the number of rows to skip.
 - **Rows to load:** Enter or use the up and down arrows to select the number of rows to load.
 - **Preview size:** Enter or use the up and down arrows to select the number of rows to preview.

To remove the options selected and the data preview, click **Clear**.

After selecting the required options, click **Apply**, and then click **Next**.

4. In Data mapping, match the data in the file to the appropriate columns in the target table. By default, the matching is done using column name.

Figure 3-9 Data Mapping

Upload Data Into HR.EMPLOYEES

1 Data preview 2 Data mapping 3 Review

File: employees1.csv **Show/Hide options**

Match columns by: Name Preview 3

Source column	Target column	Format	Row 1	Row 2	Row 3
Employee No	EMPLOYEE_ID		100	101	102
Name	- Select -		Steven	Rob	Kristy
Email	EMAIL		s@gmail.com	r@gmail.com	k@gmail.com
Phone	- Select -		9876541244	8765412345	1234567898
Hire Date	HIRE_DATE		17-06-1987	18-08-2000	18-02-2004
Last Name	LAST_NAME		CC	EE	FF

Format dropdown for Hire Date: YYYY-MM-DD, YY-MM-DD, YYYY-MON-DD, YY-MON-DD, DD/MM/RR, DD-MON-RR, RRRR-MON-DD

Buttons: Back, Next, Finish, Cancel

To modify, click **Show/Hide options** . In Match columns by:

- Select **Name** to match columns based on the name of the column in the target table.
- Select **Position** if you want to match columns based on the position of the column in the target table.
- Select **None** to remove the current selections and to select the target column for each source column from the drop-down list.

Note

Based on the data in the file, attempts are made to automatically retrieve the correct format mask of date-based columns. If this is incorrect, you can change the suggested format by entering it directly into the target cell.

If there are any issues to resolve, you see a notification such as **1 pending actions** on the top right of the dialog.

Click **Next**.

5. A summary of the previous screens is displayed. Click **Finish**.

The data will start uploading to the target table. After it is completed, an entry is added to the Log with the status of the operation. To view the Log, click the timestamp notification at the bottom of the page. If the operation is successful, a Data Import Completed notification is displayed.

6. For a detailed summary of the upload process, right-click the table in the Navigator tab, select **Data Loading**, and then select **History**. A summary of the data loaded is displayed in the Data Loading History dialog.

If any data failed to load, you can view the number of rows in the Failed Rows column. Click the column to open a dialog showing the failed rows.

In the Data Loading History dialog, you can also search for files loaded by schema name, table name, or file name. To remove the loaded files, click Remove all history .

You can also load data from a file to an existing table using the steps in [Loading Data from Multiple Local Files into Multiple Tables](#).

3.2.4.3 Loading Data from Multiple Local Files into Multiple Tables

To load data into multiple tables from multiple files:


1. In the SQL page, on the top right, click **Data Load**.

The Data Loading slider appears.

2. Click **Add File** to browse for one or more files and add them concurrently. Alternatively, you can drag and drop the files into the Data Loading slider.

The added files are displayed as cards.

3. To open a file, click the name of the file or click  at the top right corner of the card and select **Details**.



In the Details page, a preview of the file is displayed. Click **Next**. To go to the previous page, click **All Files**.

Note

A preview of the entire file is shown the first time the file is loaded. Subsequently if any changes are made to the settings, the preview is set at a maximum of 10 rows. For a description of the settings, see step 2 in [Loading Data from a Local File to a New Table](#).

4. In the Target Details page, the **Actions** field contains the following two options:
 - **Create New Table** to create a table. By default, this option is selected and the **Table Name** field is prefilled.
 - **Append to Existing Table** to add the file to an existing table. The data loader attempts to match the file name to an existing table name. Expand **Table Rows Preview** to preview the existing table.

For a description of the mapping options, see step 3 in [Loading Data from a Local File to a New Table](#) and step 4 in [Loading Data from a Local File to an Existing Table](#).

Click **All Files** to return to the initial Data Loading page.

5. To load data into a table, at the top right of the card, click  and select **Run**.

To load data into multiple tables simultaneously, click  **Run All**.

The files that are successfully uploaded are displayed with a green colour icon and a check mark and the status **UPLOADED**.

Some other possible statuses are:

- **UPLOADED WITH WARNING** in yellow.
- **NOT UPLOADED BECAUSE OF ERRORS** in red.
- **UPLOADED WITH ERRORS** where you can see the number of failed rows.
- **CONTAINS WARNING(S)** indicates that something in the settings is wrong. This file will not be loaded until the warnings are resolved.

Data Loading	
UPLOADED	Target table: PDBDBA.EMPLOYEES2
 UPLOADED	export.csv 1 row New table: PDBDBA.EXPORT_4
 UPLOADE...	numeric types 1.csv 5 total row(s), 1 row(s) failed New table: PDBDBA.NUMERIC_TYPES_1_2
 NOT UPL...	employees.csv 127 rows New table: ANONYMOUS.EMPLOYEES_13
 PENDING	export.csv 1 row New table: PDBDBA.EXPORT_5
 CONTAIN...	custom_char.tsv 1 row Target table: PDBDBA.CUSTOM_CHAR

6. After a file is loaded, go to the Details slider to see the third step named **Results**.

The Results step shows the total number of rows loaded along with the number of failed rows.

To view all previous files that have been loaded along with their loading statuses, click **Open History**. You can filter by schema and table name.

3.3 The Data Modeler Page

The Data Modeler page provides an integrated version of Oracle SQL Developer Data Modeler with basic reporting features. You can create diagrams from existing schemas, retrieve data dictionary information, generate DDL statements, and export diagrams.

To navigate to the Data Modeler page, do either of the following:

- In the Launchpad page, select the **Development** tab and click **Data Modeler**.
- Click **Selector**  to display the navigation menu. Under Development, select **Data Modeler**.

The Data Modeler page consists of the left pane for navigating objects and diagrams, the editor pane for working with relational diagrams, and the right pane for viewing the properties of the object selected in the diagram. These panes are described in the following sections:

- [Navigating Diagrams and Objects](#)
- [About the Data Modeler Editor](#)

Related Topics

- *Oracle SQL Developer Data Modeler User's Guide*

3.3.1 Navigating Diagrams and Objects

The **Diagrams** tab lists the Data Modeler diagrams that have been saved.

When you right-click a diagram, you have options to open, save, delete, and view properties. [Figure 3-10](#) shows a thumbnail preview that is displayed when you click the name of a diagram in the left pane.

Figure 3-10 Diagram Preview in Data Modeler



- To create a new diagram, in the Diagrams tab, click New Diagram .
- To open an existing diagram, in the Diagrams tab, right-click the diagram and click **Open**.

Note


Diagrams are stored in the database, in the Database Actions user schema. When you use Data Modeler for the first time, a OSDDMW_DIAGRAMS table is created in your schema, to store the diagrams.

Oracle SQL Developer Data Modeler can import from or export to the OSDDMW_DIAGRAMS table if a connection is provided. For more information, see [Sharing Diagrams with SQL Developer Data Modeler](#).

The **Navigator** tab lists the objects that are available for each schema. You can select the schema and object type from the drop-down lists. For a particular schema, if a tables object is selected, the left pane displays all the tables that belong to the particular schema. You can expand a table to view its columns. [Figure 3-11](#) indicates the important elements in the Navigator tab.

In a schema, you can search for objects in the Navigator tab by entering a few characters. You can also search the contents of a saved diagram. The search functionality is not case-sensitive, retrieves all matching entries, and does not require the use of wildcard characters.

Figure 3-11 Navigator Tab



The context menu for a table or view consists of:

Add Object to Diagram: Adds the selected object to the selected diagram. Alternatively, you can drag and drop an object into a selected diagram.

Add Object with dependencies to Diagram: Adds the selected object along with parent and child tables related to the object to the selected diagram.

Add Object as Star Schema to Diagram: Adds the selected object to the diagram and searches the data dictionary for foreign keys and implied foreign keys related to the object. The related tables or views are added to the diagram and the star schema layout is applied. See [Implied Foreign Keys](#)

Edit, Use as Template: Opens the Table Properties Dialog. Use to edit an existing object (such as a table or view) for a specific schema, or create a new object by using an existing one for the initial content. See [The Table Properties Dialog](#) and [The View Properties Dialog](#).

Add/Edit Sequence: Opens the Sequence Properties Dialog. Use to create or edit a sequence for a selected schema. See [The Sequence Properties Dialog](#).

Add Index: Opens the Index Properties Dialog. Use to create an index for a table. See [The Index Properties Dialog](#).

Add Synonym: Opens the Synonym Properties Dialog. Use to create a synonym for a table or view in a selected schema. See [The Synonym Properties Dialog](#).

3.3.2 About the Data Modeler Editor


You can create and work with relational diagrams in the editor pane.

When an object is selected in the diagram, you can inspect the properties of the object in the right pane. For a table, the properties displayed are Columns, Primary Key, Unique Constraints, Indexes, Foreign Keys, Comments, and Colors. The only properties that you can edit are background and border color. Click Reset  to return to the default colors.

[Figure 3-12](#) shows the properties of an object selected in the editor pane.

You can drag the thumbnail viewer in the upper right corner to display the appropriate area in the editor pane. This is useful when you are working with long diagrams in the editor pane.

Figure 3-12 Inspect Properties of Object Selected in Editor Pane



When creating diagrams, you can do the following:

- Select an object and then drag to move it around.
- Adjust or move objects with the relationships intact.
- Add elbows to relationship lines to avoid intersecting with lines from other objects. Right-click the relationship line and drag to create the elbow. Click Remove vertex  to restore the original shape.

- Resize objects by dragging the handles that are positioned around the box.
- Select and then right-click an object for the following options:
 - **DDL Preview:** Shows the DDL statements for the object.
 - **Update:** Updates any actions performed on the object.
 - **Delete:** Deletes the object.
 - **Edit:** Opens the Table Properties Dialog for editing a table object or the View Properties Dialog for a view object.
 - **Implied Foreign Keys:** Opens the Implied Foreign Keys Dialog. Use this option to define implied foreign keys for the object. See [Implied Foreign Keys](#)
 - **View JSON Data Guides:** Opens the JSON data guide diagram for the table or view. See [About Viewing JSON Data Guides](#)

The icons on the toolbar are:


- **Save Diagram:** Saves the currently selected diagram. Diagrams are stored in a table that is created in the schema of the user.
- **Print Diagram:** Prints the selected diagram.
- **Save to SVG:** Saves the currently selected diagram to an image file in SVG format.
- **Add Objects to Diagram:** A dialog is displayed where you can select one or more objects of type tables or views from a specific schema into the selected diagram. Type * to list all the objects of an object type (tables or views) in the schema. You can also search by typing a few characters.
- **Add Note:** Adds notes to the selected diagram. Select the note to add text and to see the associated properties in the right pane, such as Font Size and Colors.
 - To enter text, expand TEXT and type information. Click **Apply**.
 - To select the background, border, or text color for the note, click the box and select the required color. To revert to the default color, click Reset .
- **Delete:** Deletes the selected object or objects from the diagram. To select multiple objects, press the CTRL key and select the objects.
- **Zoom In:** Displays more detail, and potentially fewer objects, in the currently selected diagram.
- **Zoom Out:** Displays less detail, and potentially more objects, in the currently selected diagram.
- **Fit Screen:** Makes all relevant objects fit in the window for the currently selected diagram, adjusting the sizes of shapes and text labels as needed.
- **Actual Size:** Adjusts the shapes and text labels in the currently selected diagram to the default sizes.
- **Layout: Auto Layout:** Rearranges the objects in the diagram to a layout that may be more meaningful and attractive. If you do not like the rearrangement, you can restore the previous layout by clicking Undo.
- **Layout: Star Layout:** Rearranges the objects in the diagram to a star schema layout, where the fact table is in the center and the associated dimension tables surround the fact table.
- **DDL Preview:** Shows the DDL statements for the object. You have the option to save or send the DDL statements to the worksheet. To choose what you want to display, click **Options**.

- **DDL Preview for Current Schema:** Shows the DDL statements for the current schema. You can send the DDL statements to the worksheet.
- **Diagram Report:** Generates a data dictionary report of everything in the diagram.
- **Schema Report:** Generates a data dictionary report of everything in the schema.
- **Help:** Displays the help for the Data Modeling editor.

3.3.2.1 About Viewing JSON Data Guides

A JSON data guide represents the JSON schema for documents in one column with JSON content. A table can have more than one column with JSON content.

The View JSON Data Guides option is available in the context menu for a table or view in a Data Modeler diagram.



This is applicable for tables, views and external tables that have columns with JSON content. Columns with JSON content are identified and data guides are retrieved for each such column. This process is quicker with the existence of a JSON search index.

The JSON schema (JSON Data Guide) is presented visually like an entity-relationship diagram. Arrays are presented as one-to-many relationships, contained objects as one-to-one relationships, and "oneOf" constructs as a box that surrounds possible choices. There is a column selector at the top right part of the page enabling you to select a column with JSON content for diagram presentation.

The following example is a JSON schema and its representation in a data guide diagram.

```
{
  "type" : "object",
  "properties" :
  {
    "User" :
    {
      "type" : "string",
      "o:length" : 8,
      "o:preferred_column_name" : "DATA$User"
    },
    "PONumber" :
  }
}
```

```

{
  "type" : "number",
  "o:length" : 4,
  "o:preferred_column_name" : "DATA$PONumber"
},
"LineItems" :
{
  "type" : "array",
  "o:length" : 1024,
  "o:preferred_column_name" : "DATA$LineItems",
  "items" :
  {
    "properties" :
    {
      "Part" :
      {
        "type" : "object",
        "o:length" : 256,
        "o:preferred_column_name" : "DATA$Part",
        "properties" :
        {
          "UPCCode" :
          {
            "type" : "number",
            "o:length" : 16,
            "o:preferred_column_name" : "DATA$UPCCode"
          },
          "UnitPrice" :
          {
            "type" : "number",
            "o:length" : 8,
            "o:preferred_column_name" : "DATA$UnitPrice"
          },
          "Description" :
          {
            "type" : "string",
            "o:length" : 128,
            "o:preferred_column_name" : "DATA$Description"
          }
        }
      }
    },
    "Quantity" :
    {
      "type" : "number",
      "o:length" : 4,
      "o:preferred_column_name" : "DATA$Quantity"
    },
    "ItemNumber" :
    {
      "type" : "number",
      "o:length" : 1,
      "o:preferred_column_name" : "DATA$ItemNumber"
    }
  }
},

```

```

"Reference" :
{
  "type" : "string",
  "o:length" : 32,
  "o:preferred_column_name" : "DATA$Reference"
},
"Requestor" :
{
  "type" : "string",
  "o:length" : 16,
  "o:preferred_column_name" : "DATA$Requestor"
},
"CostCenter" :
{
  "type" : "string",
  "o:length" : 4,
  "o:preferred_column_name" : "DATA$CostCenter"
},
"ShippingInstructions" :
{
  "type" : "object",
  "o:length" : 256,
  "o:preferred_column_name" : "DATA$ShippingInstructions",
  "properties" :
  {
    "name" :
    {
      "type" : "string",
      "o:length" : 16,
      "o:preferred_column_name" : "DATA$name"
    },
    "Phone" :
    {
      "type" : "array",
      "o:length" : 64,
      "o:preferred_column_name" : "DATA$Phone",
      "items" :
      {
        "properties" :
        {
          "type" :
          {
            "type" : "string",
            "o:length" : 8,
            "o:preferred_column_name" : "DATA$type"
          },
          "number" :
          {
            "type" : "string",
            "o:length" : 16,
            "o:preferred_column_name" : "DATA$number"
          }
        }
      }
    }
  }
},
"Address" :

```

```


    {
      "oneOf" :
      [
        {
          "type" : "string",
          "o:length" : 1,
          "o:preferred_column_name" : "DATA$Address"
        },
        {
          "type" : "object",
          "o:length" : 128,
          "o:preferred_column_name" : "DATA$Address_1",
          "properties" :
          {
            "city" :
            {
              "type" : "string",
              "o:length" : 32,
              "o:preferred_column_name" : "DATA$city"
            },
            "state" :
            {
              "type" : "string",
              "o:length" : 2,
              "o:preferred_column_name" : "DATA$state"
            },
            "street" :
            {
              "type" : "string",
              "o:length" : 32,
              "o:preferred_column_name" : "DATA$street"
            },
            "country" :
            {
              "type" : "string",
              "o:length" : 32,
              "o:preferred_column_name" : "DATA$country"
            },
            "zipCode" :
            {
              "type" : "number",
              "o:length" : 8,
              "o:preferred_column_name" : "DATA$zipCode"
            }
          }
        }
      ]
    },
    "Special Instructions" :
    {
      "type" : "string",
      "o:length" : 16,
      "o:preferred_column_name" : "DATA$SpecialInstructions"
    }
  }

```

```

    }
  }
}

```

Figure 3-13 JSON Data Guide Diagram**See Also**

[JSON Data Guide in Oracle Database JSON Developer's Guide](#)

3.4 Creating and Editing Database Objects

You can create and edit objects using Create and Edit Object wizards available from the Navigator tab in the SQL and Data Modeler pages.

The wizards for creating and editing various object types are described in the following sections:

- [The Table Properties Dialog](#)
- [The Index Properties Dialog](#)
- [The Sequence Properties Dialog](#)
- [The View Properties Dialog](#)
- [The Synonym Properties Dialog](#)
- [Implied Foreign Keys](#)
- [The Materialized View Log Properties Dialog](#)


3.4.1 The Table Properties Dialog

The Table Properties Dialog is displayed when you create a table, edit an existing table, or create a table using an existing one as a template.

You can open the Table Properties dialog from the Navigator tab in SQL or Data Modeler.

To create a table for a specific schema, in the Navigator tab, select **Tables** from the object type drop-down list, click **Object submenu** **...**, and select **Create Object**.


Figure 3-14 Create Object



To create a table from an existing one for a specific schema, right-click the table object in the Navigator tab, and select **Use as Template** as shown in [Figure 3-15](#).

To edit a table for a specific schema, right-click a table object in the Navigator tab, and select **Edit**.

Figure 3-15 Edit Option for Existing Table



The table properties are grouped in several panes.

If you are editing an existing table, you can visit the panes in any order. If you click Create before you finish creating the table, right-click the table name, select **Edit**, and continue creating the table.

Note

Editing a partitioned table is not recommended. To identify whether a table is partitioned or not, right-click the table name and select **Edit**. If the table is partitioned, a warning message will be displayed.

Schema: Database schema in which to create the table. By default, a new table is created in the existing schema or the schema that you are logged into.

Name: Name for the table.

The different panes in the dialog are described in the following sections:

- [Columns Pane](#)
- [Primary Key Pane](#)
- [Unique Keys Pane](#)
- [Indexes Pane](#)
- [Foreign Keys Pane](#)
- [Table Constraints Pane](#)
- [Comments Pane](#)
- [Storage Pane](#)
- [External Table Properties Pane](#)
- [Materialized View Pane](#)
- [DDL Pane](#)
- [Output Pane](#)

Related Topics

- *Oracle Database SQL Language Reference*

3.4.1.1 Columns Pane


Specifies properties for each column in the table.

General tab

Lists the columns available in the table.

To add a column, click **Add Column (+)**. A new row is added to the table below. Select the row and enter the details for the column.

To delete a column, select the row and click **Remove Column (-)**. To move a column up or down in the table, select it and use the up-arrow and down-arrow icons.

Figure 3-16 Add a Column

- **Name:** Name for the column.
- **Datatype:** Data type for the column.
- **Default:** If no value is specified, the default value inserted into the column when a row is inserted.
- **Default on NULL:** Applicable for Oracle Database 12c and later releases. If this option is selected, when a row is inserted into the table and the value specified for the column is NULL, the default value is inserted into the column.
- **Expression:** Expression for computing the value in the column.
- **Comments:** Optional descriptive comments about the column. Use this field to provide descriptions for the attributes.

In the table:

- **PK:** If this option is selected, the column becomes the primary key.
- **Identity Column:** If this option is selected, the column becomes an identity column. This is applicable only for Oracle Database 12c and later releases. For more details, see the Identity Column tab.

Constraints tab

Displays the Not Null and Check Constraints for a column. A check constraint requires values in a column to comply with a specified condition.

- **Not Null Constraint:** **Name:** Name for the Not Null constraint.

- **Not Null Constraint: Not Null:** If this option is selected, the column must contain data. You cannot specify no value or an explicit null value for this column when you insert a row. If this option is not checked, the column can contain either data or no data. A primary key column cannot be null.
- **Check Constraint: Name:** Name for the check constraint definition.
- **Check Constraint: Constraint:** Condition that must be met for a column to fulfill the check constraint. You can use any valid CHECK clause (without the CHECK keyword). For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify: rating >=1 and rating <= 10.
- **Enabled:** If this option is selected, the constraint is checked when data is entered or updated in the column.
- **Deferrable:** If this option is selected, you can defer checking the validity of the constraint until the end of a transaction.
- **Initially Immediate:** If this option is selected, the constraint is checked whenever you add, update, or delete data from the column.
- **Validate:** If this option is selected, the existing data is checked to see if it conforms to the constraint.

Identity Column tab

Applicable for Oracle Database 12c and later releases. The Identity Column tab lists the properties of the identity column. This tab becomes available only after the Identity Column checkbox is selected for the column in the General tab. An identity column is an autoincrement column that can be used to identify a table row. Only one identity column can be specified for a table.

- **Generate:** Always means that values cannot be explicitly included for the identity column in INSERT OR UPDATE statements, By Default means values for the identity column are generated automatically if no values are specified explicitly, By Default on Null means values are generated for the column only when a NULL value is supplied.
- **Start with:** Starting value of the sequence.
- **Increment:** Interval between successive numbers in a sequence.
- **Min value:** Lowest possible value for the sequence. The default is 1 for an ascending sequence and $-(10^{26})$ for a descending sequence.
- **Max value:** Highest possible value for the sequence. The default is 10^{27} for an ascending sequence and -1 for a descending sequence.
- **Cache and Cache size:** Cache causes sequence values to be preallocated in cache, which can improve application performance; Cache size indicates the number of sequence values preallocated in cache. No Cache causes sequence values not to be preallocated in cache.
- **Cycle:** Indicates whether the sequence "wraps around" to reuse numbers after reaching its maximum value (for an ascending sequence) or its minimum value (for a descending sequence). If cycling of values is not enabled, the sequence cannot generate more values after reaching its maximum or minimum value.
- **Order:** Indicates whether sequence numbers are generated in the order in which they are requested. If No Order is specified, sequence numbers are not guaranteed to be in the order in which they were requested.

3.4.1.2 Primary Key Pane

Specifies the primary key for the table.

The primary key is the column, or set of columns, that uniquely identifies each row in the table. If the Primary Key checkbox is selected for a column in the General tab, the corresponding fields are automatically populated in the Primary Key pane. You can make changes to the properties as required.

An index is automatically created on the primary key.

- **Name:** Name of the constraint to be associated with the primary key definition.
- **Enabled:** If this option is checked, the primary key constraint is enforced: that is, the data in the primary key column (or set of columns) must be unique and not null.
- **Index:** Name of the index to which the primary key refers.
- **Tablespace:** Name of the tablespace associated with the index.
- **Available Columns:** Lists the columns that are available to be added to the primary key definition. You can select multiple attributes, if required, for the primary key.
- **Selected Columns:** Lists the columns that are included in the primary key definition.

To add a column to the primary key definition, select it in Available Columns and click the Add (>) icon; to remove a column from the primary key definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the primary key definition, select it in Selected Columns and use the arrow buttons.

3.4.1.3 Unique Keys Pane

Specifies one or more unique constraints for the table.

A unique constraint specifies a column, or set of columns, whose data values must be unique: each data value must not be null, and it must not be the same as any other value in the column.

To add a unique constraint, click the Add button; to delete a unique constraint, select it and click the Remove button.

- **Name:** Name of the unique constraint.
- **Enabled:** If this option is selected, the unique constraint is enforced.
- **Rely:** If this option is selected, the constraint in NOVALIDATE mode is taken into account during query rewrite.
- **Deferrable:** If this option is selected, in subsequent transactions, constraint checking can be deferred until the end of the transaction using the SET CONSTRAINT(S) statement.
- **Initially Immediate:** If this option is selected, the constraint is checked at the end of each subsequent SQL statement.
- **Validate:** If the option is selected, the existing data is checked to see if it conforms to the constraint.
- **Index:** Name of the index to which the unique key refers.
- **Tablespace:** Name of the tablespace associated with the index.
- **Available Columns:** Lists the columns that are available to be added to the unique constraint definition.

- **Selected Columns:** Lists the columns that are included in the unique constraint definition.

To add a column to the unique constraint definition, select it in Available Columns and click the Add (>) icon; to remove a column from the unique constraint definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the unique constraint definition, select it in Selected Columns and use the arrow buttons.

3.4.1.4 Indexes Pane

Lists the indexes defined for the table.

To add an index, click Add Index (+); to delete an index, select it and click Remove Index (-).

- **Name:** Name of the index.
- **Type:** The type of Oracle index. Non-unique means that the index can contain multiple identical values; Unique means that no duplicate values are permitted; Bitmap stores rowids associated with a key value as a bitmap.
- **Tablespace:** Name of the tablespace for the index.
- **Expression:** A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.
- **Available Columns and Selected Columns:** Columns selected for the index. To select a column, click the column in the Available Columns box, and then click the Add Selected Columns icon to move it to the Selected Columns box.

3.4.1.5 Foreign Keys Pane

Specifies one or more foreign keys for the table.

A foreign key specifies a column ("local column"), whose data values match values in the primary key or unique constraint of another table.

- **Name:** Name of the foreign key definition.
- **Enabled:** If this option is checked, the foreign key is enforced.
- **Rely, Deferrable, Initially Immediate, Validate:** See the description of these fields in the Unique Keys pane.
- **Referenced Constraint: Schema:** Name of the schema containing the table with the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: Table:** Name of the table with the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: Constraint:** Name of the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: On Delete:** Action to take automatically when a row in the referenced table is deleted and rows with that value exist in the table containing this foreign key: NO ACTION (shown by a crossing line in diagrams) performs no action on these rows; CASCADE (shown by an "X") deletes these rows; SET NULL (shown by a small circle) sets null all columns in those rows that can be set to a null value.
- **Associations: Local Column:** Lists the column in the currently selected (local) table that is included in the foreign key definition. For each referenced column in the foreign key definition, select the name of a column in the edited table.

- **Associations: Referenced Column:** For each local column, identifies the column in the other (foreign) table that must have a value matching the value in the local column.

3.4.1.6 Table Constraints Pane

Specifies one or more check constraints for the table.

A check constraint specifies a condition that must be met when a row is inserted into the table or when an existing row is modified.

- **Name:** Name of the check constraint definition.
- **Check Condition:** Condition that must be met for a row to fulfil the check constraint. You can use any valid CHECK clause (without the CHECK keyword). For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify rating >=1 and rating <= 10.
- **Enabled:** If this option is checked, the check constraint is enforced.

3.4.1.7 Comments Pane

Enter descriptive comments in this pane. This is optional.

3.4.1.8 Storage Pane

Enables you to specify storage options for the table.

When you create or edit a table or an index, you can override the default storage options.

- **Organization:** Specifies that the table is stored and organized with (Index) or without an index (Heap) or as an external table (External).
- **Tablespace:** Name of the tablespace for the table or index.
- **Logging:** ON means that the table creation and any subsequent INSERT operations against the table are logged in the redo log file. OFF means that these operations are not logged in the redo log file.
- **Row Archival:** YES enables in-database archiving, which allows you to archive rows within the table by marking them as invisible.

3.4.1.9 External Table Properties Pane

Specifies options for an external table.

An external table is a read-only table whose metadata is stored in the database but whose data is stored outside the database.

External Table

- **Access Driver Type:** Specifies the type of external table.
 - ORACLE_LOADER: Extracts data from text data files. This is the default access driver, which loads data from external tables to internal tables.
 - ORACLE_DATAPUMP: Extracts data from binary dump files. This access driver can perform both loads and unloads.
 - ORACLE_BIGDATA: Extracts data from Oracle Big Data Appliance.
 - ORACLE_HDFS: Extracts data stored in a Hadoop Distributed File System (HDFS).

- ORACLE_HIVE: Extracts data stored in Apache HIVE.
- **Default Directory:** Specifies the default directory to use for all input and output files that do not explicitly name a directory object. The location is specified with a directory object, not a directory path.
- **Access Params:** Assigns values to the parameters of the specific access driver for the external table. Access parameters are optional.
 - OPAQUE_FORMAT_SPEC: The opaque_format_spec specifies all access parameters for the ORACLE_LOADER, ORACLE_DATAPUMP, ORACLE_HDFS, and ORACLE_HIVE access drivers. For descriptions of the access parameters, see *Oracle Database Utilities*. Field names specified in the opaque_format_spec must match columns in the table definition, else Oracle Database ignores them.
 - USING CLOB: Enables you to derive the parameters and their values through a subquery. The subquery cannot contain any set operators or an ORDER BY clause. It must return one row containing a single item of data type CLOB.
- **Reject Limit:** The number of conversion errors that can occur during a query of the external data before an Oracle Database error is returned and the query is aborted.
- **Project Column:** Determines how the access driver validates the rows of an external table in subsequent queries.
 - ALL: Processes all column values, regardless of which columns are selected, and validates only those rows with fully valid column entries. If any column value raises an error, such as a data type conversion error, the row is rejected even if that column was not referenced in the select list of the query.
 - REFERENCED: Processes only those columns in the select list of the query.

The ALL setting guarantees consistent result sets. The REFERENCED setting can result in different numbers of rows returned, depending on the columns referenced in subsequent queries, but is faster than the ALL setting. If a subsequent query selects all columns of the external table, then the settings behave identically.

- **Location:** Specifies the data files for the external table. Use the Add (+) icon to add each location specification.
 - For ORACLE_LOADER and ORACLE_DATAPUMP, the files are named in the form `directory:file`. The directory portion is optional. If it is missing, then the default directory is used as the directory for the file. If you are using the ORACLE_LOADER access driver, then you can use wildcards in the file name. An asterisk (*) signifies multiple characters and a question mark (?) signifies a single character.
 - For ORACLE_HDFS, LOCATION is a list of Uniform Resource Identifiers (URIs) for a directory or for a file. There is no directory object associated with a URI.
 - For ORACLE_HIVE, LOCATION is not used. Instead, the Hadoop HCatalog table is read to obtain information about the location of the data source (which could be a file or another database).

Opaque Format Spec

Specifies all access parameters for the ORACLE_LOADER, ORACLE_DATAPUMP, ORACLE_HDFS, and ORACLE_HIVE access drivers.

For example:

```
RECORDS DELIMITED BY NEWLINE CHARACTERSET US7ASCII
TERRITORY AMERICA
BADFILE log_file_dir:'ext_1v3.bad'
```

```

LOGFILE log_file_dir:'ext_1v3.log'
FIELDS TERMINATED BY "|" OPTIONALLY ENCLOSED BY '^' LDRTRIM
( PROD_ID,
  CUST_ID ,
  TIME_ID DATE(10) "YYYY-MM-DD",
  CHANNEL_ID ,
  PROMO_ID ,
  QUANTITY_SOLD ,
  AMOUNT_SOLD ,
  UNIT_COST ,
  UNIT_PRICE
)

```

and the full statement:

```

CREATE TABLE SH.SALES_TRANSACTIONS_EXT
(
  PROD_ID NUMBER ,
  CUST_ID NUMBER ,
  TIME_ID DATE ,
  CHANNEL_ID NUMBER ,
  PROMO_ID NUMBER ,
  QUANTITY_SOLD NUMBER ,
  AMOUNT_SOLD NUMBER (10,2) ,
  UNIT_COST NUMBER (10,2) ,
  UNIT_PRICE NUMBER (10,2)
)
ORGANIZATION EXTERNAL
(
  TYPE ORACLE_LOADER
  DEFAULT DIRECTORY DATA_FILE_DIR
  ACCESS PARAMETERS
  (
    RECORDS DELIMITED BY NEWLINE CHARACTERSET US7ASCII
  TERRITORY AMERICA
  BADFILE log_file_dir:'ext_1v3.bad'
  LOGFILE log_file_dir:'ext_1v3.log'
  FIELDS TERMINATED BY "|" OPTIONALLY ENCLOSED BY '^' LDRTRIM
  ( PROD_ID ,
    CUST_ID ,
    TIME_ID DATE(10) "YYYY-MM-DD",
    CHANNEL_ID ,
    PROMO_ID ,
    QUANTITY_SOLD ,
    AMOUNT_SOLD ,
    UNIT_COST ,
    UNIT_PRICE
  )
)
LOCATION ( "DATA_FILE_DIR":'sale1v3.dat')
)
REJECT LIMIT 100
;

```

CLOB Subquery

Type or copy and paste the query.

Note

For more information about the external table fields, see *Oracle Database Utilities* and *Oracle Database SQL Language Reference*

3.4.1.10 Materialized View Pane

Specifies options for a materialized view.

Query: Contains the SQL code for the query part of the view definition. Type or copy and paste the query.

General

- **On Pre-built Table:** If **Yes**, an existing table is registered as a preinitialized materialized view. This option is particularly useful for registering large materialized views in a data warehousing environment. The table must have the same name and be in the same schema as the resulting materialized view, and the table should reflect the materialization of a subquery.
- **Reduced Precision:** **Yes** authorizes the loss of precision that will result if the precision of the table or materialized view columns do not exactly match the precision returned by the subquery. If **No**, the precision of the table or materialized view columns must exactly match the precision returned by the subquery, or the create operation will fail.
- **For Update:** Select **Yes** to allow a subquery, primary key, object, or rowid materialized view to be updated. When used in conjunction with Advanced Replication, these updates will be propagated to the master.
- **Real Time MV:** Select **Yes** to create a real-time materialized view or a regular view. A real-time materialized view provides fresh data to user queries even when the materialized view is not in sync with its base tables due to data changes. Instead of modifying the materialized view, the optimizer writes a query that combines the existing rows in the materialized view with changes recorded in log files (either materialized view logs or the direct loader logs). This is called on-query computation.
- **Query Rewrite:** If **Enable**, the materialized view is enabled for query rewrite, which transforms a user request written in terms of master tables into a semantically equivalent request that includes one or more materialized views.
- **Build:** Specifies when to populate the materialized view. **Immediate** indicates that the materialized view is to be populated immediately. **Deferred** indicates that the materialized view is to be populated by the next refresh operation. If you specify **Deferred**, the first (deferred) refresh must always be a complete refresh; until then, the materialized view has a staleness value of unusable, so it cannot be used for query rewrite.
- **Use Index:** If **Yes**, a default index is created and used to speed up incremental (fast) refresh of the materialized view. If **No**, this default index is not created. (For example, you might choose to suppress the index creation now and to create such an index explicitly later.)
- **Index Tablespace:** Specifies the tablespace in which the materialized view is to be created. If a tablespace is not selected, the materialized view is created in the default tablespace of the schema containing the materialized view.

- **Cache:** If **Yes**, the blocks retrieved for this table are placed at the most recently used end of the least recently used (LRU) list in the buffer cache when a full table scan is performed. This setting is useful for small lookup tables. If **No**, the blocks are placed at the least recently used end of the LRU list.

Refresh Clause

- **Refresh:** Select Yes to enable refresh operations.
- **Refresh Type:** The method of refresh operation to be performed:
 - Complete Refresh: Executes the defining query of the materialized view, even if a fast refresh is possible.
 - Fast Refresh: Uses the incremental refresh method, which performs the refresh according to the changes that have occurred to the master tables. The changes for conventional DML changes are stored in the materialized view log associated with the master table. The changes for direct-path INSERT operations are stored in the direct loader log.
 - Force Refresh: Performs a fast refresh if one is possible; otherwise, performs a complete refresh.
- **Action:** The type of refresh operation to be performed:
 - On Demand: Performs a refresh when one of the DBMS_MVIEW refresh procedures are called.
 - On Commit: Performs a fast refresh whenever the database commits a transaction that operates on a master table of the materialized view. This may increase the time taken to complete the commit, because the database performs the refresh operation as part of the commit process.
 - Specify: Performs refresh operations according to what you specify in the Start on and Next fields.
- **Start Date:** Starting date and time for the first automatic refresh operation. Must be in the future.
- **Next Date:** Time for the next automatic refresh operation. The interval between the Start on and Next times establishes the interval for subsequent automatic refresh operations. If you do not specify a value, the refresh operation is performed only once at the time specified for Start on.
- **With:** Refresh type, which determines the type of materialized view:
 - Primary Key: Creates a primary key materialized view, which allows materialized view master tables to be reorganized without affecting the eligibility of the materialized view for fast refresh.
 - Row ID: Creates a rowid materialized view, which is useful if the materialized view does not include all primary key columns of the master tables.
- **Default Storage:** If Yes, DEFAULT specifies that Oracle Database will choose automatically which rollback segment to use. If you specify DEFAULT, you cannot specify the rollback_segment. DEFAULT is most useful when modifying, rather than creating, a materialized view.
- **Storage Type:** MASTER specifies the remote rollback segment to be used at the remote master site for the individual materialized view. LOCAL specifies the remote rollback segment to be used for the local refresh group that contains the materialized view. This is the default.
- **Rollback Segment:** Enter the name of the rollback segment.

- **Using Constraint:** If this option is checked, more rewrite alternatives can be used during the refresh operation, resulting in more efficient refresh execution. The behavior of this option is affected by whether you select Enforced or Trusted.
 - **Enforced:** Causes only enforced constraints to be used during the refresh operation.
 - **Trusted:** Enables the use of dimension and constraint information that has been declared trustworthy by the database administrator but that has not been validated by the database. If the dimension and constraint information is valid, performance may improve. However, if this information is invalid, then the refresh procedure may corrupt the materialized view even though it returns a success status.

3.4.1.11 DDL Pane

You can review and save the SQL statements that are generated when creating or editing the object. If you want to make any changes, go back to the relevant panes and make the changes there.

- For a new table, click **CREATE** to view the generated DDL statements.
- When you edit table properties, click **UPDATE** to view the generated ALTER statements. For a new table, the UPDATE tab will not be available.

When you are finished, click **Apply**.

3.4.1.12 Output Pane

Displays the results of the DDL commands. If there are any errors, go to the appropriate pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

3.4.2 The Index Properties Dialog

The Index Properties dialog box is displayed when you create or edit an index.

To create an index for a selected schema, in SQL, in the Navigator tab, select **Indexes** from the object type drop-down list, click Object submenu *******, and select **Create Object**.

To edit an index for a selected schema, right-click a table object in the Navigator tab, and select **Edit**.

Definition pane

- **Schema:** Database schema that owns the table associated with the index.
- **Table:** Name of the table associated with the index.
- **Schema:** Database in which to create the index.
- **Tablespace:** Tablespace for the index.
- **Name:** Name of the index.
- **Type:** The type of Oracle index.
 - **Non-unique** means that the index can contain multiple identical values.
 - **Unique** means that no duplicate values are permitted.
 - **Bitmap** stores rowids associated with a key value as a bitmap.

- **Expression:** A column name or column expression. A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.
- **Available Columns:** Columns available in the table.
- **Selected Columns:** Columns selected for the index. Click Add Selected Columns > to move columns from the Available Columns list.
- **Order:** **ASC** for an ascending index (index values sorted in ascending order); **DESC** for a descending index (index values sorted in descending order).

DDL pane

You can review and save the SQL statements that are generated when creating or editing the index. If you want to make any changes, go back to the Definition pane and make the changes there.

- For a new index, click **CREATE** to view the generated DDL statements.
- When you edit index properties, click **UPDATE** to view the generated ALTER statements. For a new index, the UPDATE tab will not be available.

When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the Definition pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

3.4.3 The Sequence Properties Dialog

The Sequence Properties Dialog is displayed when you create or edit a sequence.

To create a sequence for a selected schema, in SQL, in the Navigator tab, select **Sequences** from the object type drop-down list, click Object submenu *******, and select **Create Object**.

To edit a sequence for a selected schema, right-click a sequence object in the Navigator tab and select **Edit**.

A sequence is an object from which multiple users may generate unique integers. You can use sequences to automatically generate primary key values.

Properties pane

- **Schema:** Database schema in which to create the sequence.
- **Name:** Name of the sequence.
- **Start with:** Starting value of the sequence.
- **Increment:** Interval between successive numbers in a sequence.
- **Min value:** Lowest possible value for the sequence. The default is 1 for an ascending sequence and $-(10^{26})$ for a descending sequence.
- **Max value:** Highest possible value for the sequence. The default is 10^{27} for an ascending sequence and -1 for a descending sequence.
- **Cache and Cache size:** Cache causes sequence values to be preallocated in cache, which can improve application performance; Cache size indicates the number of sequence values preallocated in cache. No Cache causes sequence values not to be preallocated in cache.

- **Cycle:** Indicates whether the sequence "wraps around" to reuse numbers after reaching its maximum value (for an ascending sequence) or its minimum value (for a descending sequence). If cycling of values is not enabled, the sequence cannot generate more values after reaching its maximum or minimum value.
- **Order:** Indicates whether sequence numbers are generated in the order in which they are requested. If **No Order** is specified, sequence numbers are not guaranteed to be in the order in which they were requested.

DDL pane

You can review and save the SQL statements that are generated when creating or editing the sequence. If you want to make any changes, go back to the Properties pane and make the changes there.

- For a new sequence, click **CREATE** to view the generated DDL statements.
- When you edit a sequence, click **UPDATE** to view the generated ALTER statements. For a new sequence, the UPDATE tab will not be available.

When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the Properties pane, fix the errors, and run the commands again. You have save to a text file or clear the output.

3.4.4 The View Properties Dialog

The View Properties Dialog is displayed when you create or edit a view.

You can open the View Properties Dialog from the Navigator tab in SQL or Data Modeler.

To create a view for a selected schema, in SQL, in the Navigator tab, select **Views** from the object type drop-down list, click Object submenu ******* , and select **Create Object**.

To create a view from an existing template for a selected schema, in the Navigator tab, select the view to create from, right-click and select **Use as Template**.

To edit a view for a selected schema, right-click a view object in the Navigator pane, and select **Edit**.

Schema: Database schema in which to create the view.

Name: Name of the view.

The different panes in the dialog are described in the following sections:

SQL Query pane

Enter or copy and paste the SQL query for the view, using the SELECT and FROM keywords along with the syntax needed to retrieve the desired information. A semicolon is not required after the query.

Columns pane

Click Refresh Columns  to automatically populate the columns in this pane. You can edit the columns by selecting the required row and making changes in the Header Alias and Comments fields.

Storage pane

- **Force on Create:** Select **Yes** to create the view regardless of whether the base tables of the view or the referenced object types exist or the owner of the schema containing the view has privileges on them. These conditions must be true before any **SELECT**, **INSERT**, **UPDATE**, or **DELETE** statements can be issued against the view. If the view definition contains any constraints, **CREATE VIEW ... FORCE** fails if the base table does not exist or the referenced object type does not exist. **CREATE VIEW ... FORCE** also fails if the view definition names a constraint that does not exist.
- **Query Restriction: Read Only** prevents the view from being used to add, delete, or change data in the underlying table. **Check Option** prohibits any changes to the underlying table that would produce rows that are not included in this view.

Use the Primary Key, Unique Keys, Foreign Keys, and Comments panes to add or edit properties as required.

DDL pane

Based on the inputs provided, the DDL statements are generated. You can review and save the SQL statements. If you want to make any changes, go back to the relevant pane and make the changes there.

- For a new view, click **CREATE** to view the generated DDL statements.
- When you edit a view, click **UPDATE** to view the generated ALTER statements. For a new view, the UPDATE tab will not be available.

When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the respective pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

3.4.5 The Synonym Properties Dialog

The Synonym Properties Dialog is displayed when you create a synonym.

There are two ways of creating a synonym for a selected schema:

- In SQL, in the Navigator tab, right-click the object for which you want to create the synonym, and select **Add Synonym**. In this case, the only fields that you can edit in the Properties pane are Public and Synonym Name. The values of the remaining fields are predetermined by the object selected.
- In SQL, in the Navigator tab, select the object type as **Synonyms** or **Public Synonyms** from the drop-down list. Click Object submenu *******, and select **Create Object**. All the fields in the Properties dialog are available for edit.

The different panes in the dialog are described in the following sections:

Properties pane

- **Public:** If this option is checked, the synonym is accessible to all users. However, each user must have appropriate privileges on the underlying object to use the synonym. If this option is not checked, the synonym is a private synonym, and is accessible only within its schema.
- **Synonym Schema:** Database schema in which to create the synonym.
- **Synonym Name:** Name of the synonym. A private synonym must be unique within its schema; a public synonym must be unique within the database.
- **Object Type:** Specify the type of object to which this synonym refers.

- **Object Schema:** Schema containing the object or name to which this synonym refers.
- **DB Filter:** After selecting the Object Type and Object Schema, the list of objects of the selected type may be very long. To filter the object names, enter the search entry and click Refresh . The Object Name field is auto-filled with appropriate object names in the drop-down list.
- **Object Name:** Select the name of the object to which this synonym refers.
- **DB Link:** Enter a complete or partial database link to create a synonym for a schema object on a remote database where the object is located. If you specify DB Link and omit schema, then the synonym refers to an object in the schema specified by the database link. Oracle recommends that you specify the schema containing the object in the remote database. If you omit DB Link, then Oracle Database assumes the object is located on the local database.

DDL pane

Based on the inputs provided, the DDL statements are generated. You can review and save the SQL statements. If you want to make any changes, go back to the relevant pane and make the changes there.

- For a new view, click **CREATE** to view the generated DDL statements.
- When you edit a view, click **UPDATE** to view the generated ALTER statements. For a new view, the UPDATE tab will not be available.

When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the respective pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

3.4.6 The Materialized View Log Properties Dialog

The Materialized View Log Properties dialog is displayed when you create or edit a materialized view log, which is a table associated with the master table of a materialized view.

To create a materialized view log for a selected schema, in SQL, in the Navigator tab, select **Materialized View Logs** from the object type drop-down list, click Object submenu , and select **Create Object**.

To edit, right-click a materialized view log object in the Navigator pane and select **Edit**.

Schema: Database schema in which to create the materialized view log.

Table: Name of the master table of the materialized view to be associated with this materialized view log.

Properties tab

- **Row ID Logged:** **Yes** indicates that the rowid of all rows changed should be recorded in the materialized view log. **No** indicates that the rowid of all rows changed should not be recorded in the materialized view log.
- **PK Logged:** **Yes** indicates that the primary key of all rows changed should be recorded in the materialized view log; **No** indicates that the primary key of all rows changed should not be recorded in the materialized view log.
- **New values:** **Yes** saves both old and new values for update DML operations in the materialized view log; **No** disables the recording of new values in the materialized view log.

If this log is for a table on which you have a single-table materialized aggregate view, and if you want the materialized view to be eligible for fast refresh, you must specify **Yes**.

- **Object ID Logged:** For a log on an object table only: **Yes** indicates that the system-generated or user-defined object identifier of every modified row should be recorded in the materialized view log. **No** indicates that the system-generated or user-defined object identifier of every modified row should not be recorded in the materialized view log.
- **Cache:** For data that will be accessed frequently, **CACHE** specifies that the blocks retrieved for this log are placed at the most recently used end of the least recently used list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables. **NOCACHE** specifies that the blocks are placed at the least recently used end of the LRU list.
- **Parallel:** If **YES**, parallel operations will be supported for the materialized view log.
- **Sequence Logged:** **Yes** indicates that a sequence value providing additional ordering information should be recorded in the materialized view log. **No** indicates that a sequence value providing additional ordering information should not be recorded in the materialized view log. Sequence numbers (that is, **Yes** for this option) are necessary to support fast refresh after some update scenarios.
- **Commit SCN:** If this option is enabled, the database is instructed to use commit SCN data rather than timestamps.
- **Available Columns** and **Selected Columns:** Additional columns, which are non-primary-key columns referenced by subquery materialized views, to be recorded in the materialized view log. To select one or more filter columns, use the arrow buttons to move columns from Available to Selected.

Storage tab

- **Tablespace:** Tablespace in which the materialized view log is to be created
- **Logging:** **YES** or **NO**, to establish the logging characteristics for the materialized view log.
- **Buffer Mode:** Select **KEEP** to put blocks from the segment into the KEEP buffer pool. Select **RECYCLE** to put blocks from the segment into the RECYCLE pool. Select **DEFAULT** to indicate the default buffer pool.
- **Percent Free:** Specify a whole number representing the percentage of space in each data block of the database object reserved for future updates to rows of the object. The value of **PCTFREE** must be a value from 0 to 99.
- **Percent Used:** Specify a whole number representing the minimum percentage of used space that Oracle maintains for each data block of the database object. **PCTUSED** is specified as a positive integer from 0 to 99 and defaults to 40.
- **Initrans:** Specify the initial number of concurrent transaction entries allocated within each data block allocated to the database object. This value can range from 1 to 255 and defaults to 1.
- **Freelists:** In tablespaces with manual segment-space management, for objects other than tablespaces and rollback segments, specify the number of free lists for each of the free list groups for the table, partition, cluster, or index. The default and minimum value for this parameter is 1, meaning that each free list group contains one free list.
- **Freelist Groups:** In tablespaces with manual segment-space management, specify the number of groups of free lists for the database object you are creating.
- **Initial Extent:** Specify the size of the first extent of the object.
- **Next Extent:** Specify in bytes the size of the next extent to be allocated to the object.

- **Percent Increase:** In locally managed tablespaces, Oracle Database uses the value of PCTINCREASE during segment creation to determine the initial segment size and ignores this parameter during subsequent space allocation.
- **Min Extent:** In locally managed tablespaces, Oracle Database uses the value of MINEXTENTS in conjunction with PCTINCREASE, INITIAL and NEXT to determine the initial segment size.
- **Max Extent:** This storage parameter is valid only for objects in dictionary-managed tablespaces. Specify the total number of extents, including the first, that Oracle can allocate for the object.
- **Unlimited:** Select this option if you want extents to be allocated automatically as needed. Oracle recommends this setting as a way to minimize fragmentation.

Purge tab

- **Type:** In **IMMEDIATE SYNCHRONOUS**, the materialized view log is purged immediately after refresh. This is the default. In **IMMEDIATE ASYNCHRONOUS**, the materialized view log is purged in a separate Oracle Scheduler job after the refresh operation.
- **Deferred, Start With, Next, Repeat Interval:** Sets up a scheduled purge that is independent of the materialized view refresh and is initiated during CREATE or ALTER MATERIALIZED VIEW LOG statement.

Refresh tab

- **Type: Synchronous Refresh** creates a staging log that can be used for synchronous refresh. Specify the name of the staging log to be created. The staging log will be created in the schema in which the master table resides. **Fast Refresh** creates a materialized view log that can be used for fast refresh. The materialized view log will be created in the same schema in which the master table resides. This is the default.

DDL pane

Based on the inputs provided, the DDL statements are generated. You can review and save the SQL statements. If you want to make any changes, go back to the relevant pane and make the changes there.

- For a new materialized view log, click **CREATE** to view the generated DDL statements.
- When you edit a materialized view log, click **UPDATE** to view the generated ALTER statements. For a new materialized view log, the UPDATE tab will not be available.

When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the respective pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

3.4.7 Implied Foreign Keys

Implied foreign keys are dependencies that exist between tables but are not defined in the database. In a data warehouse environment, it is a common practice not to create foreign keys. However, it becomes necessary to show these dependencies for presentation or reporting purposes.

You can display implied foreign keys for objects in a star schema by defining them or by discovering them in the data dictionary. It is possible to have more than one source for implied foreign keys.

Define Implied Foreign Keys

You can define implied foreign keys in two ways:


- Using the Implied Foreign Keys dialog
- By dragging the arrow to the referenced object in the diagram

Using the Implied Foreign Keys Dialog

You can define implied foreign keys using the Implied Foreign Keys dialog in Data Modeler.

1. In a Data Modeler diagram, right-click an object (table or view) and select **Implied Foreign Keys**.

Figure 3-17 Select Implied Foreign Keys for an Object



The Implied Foreign Keys dialog is displayed.

2. In the Implied Foreign Keys dialog, click + to add an entry in the grid.
3. Select the entry in the grid to enable and enter values in the following fields:
 - **Referenced Object:** Object in the diagram that has a dependency to the source object.
 - **Local Column:** Name of the column in the source object.
 - **Referenced Column:** Name of the column in the targeted object.
 - **Discovery Sources:** Automatically prefilled, displays whether the implied foreign keys have been defined or were discovered in the data dictionary.
4. Click **OK**. The implied foreign key dependency is displayed with a dotted line on the diagram.

Figure 3-18 Dotted Line Between Two Objects




By Dragging the Arrow to the Referenced Object in the Diagram

You can also define an implied foreign key in the following way:

1. Select the source object on the diagram.
2. Click and drag the small curved arrow with a blue indicator to the referenced object. The dependency will be displayed with a dotted line on the diagram.

Figure 3-19 Drag Blue Indicator



3. Right-click the source object and enter the column names in the Implied Foreign Keys dialog.

Discover Implied Foreign Keys in the Data Dictionary

A star schema can be discovered by searching for several types of definitions in the data dictionary.

In the Navigator tab, right-click an object (table or view) and select **Add Object as Star Schema to Diagram**. The object must be a fact table. The data dictionary is then searched for joins and dependencies related to the object, such as:

- Foreign keys defined for the selected table to other tables. If implied foreign keys are later discovered for the same columns, they will not be displayed on the diagram.
- Joins used in the definitions of bitmap join indexes, materialized views with aggregates, and attribute clustering.
- Dependencies based on dimension definitions and column name matching in fact table.
- Fact and dimension definitions for Analytic views and OLAP cube and dimensions.

4

The REST Pages

The REST pages enable you to develop and protect ORDS (Oracle REST Data Services) based RESTful web services for your database.

To navigate to the REST pages, do either of the following:

- In the Launchpad page, select the **Development** tab and click **REST**.
- Click **Selector**  to display the navigation menu. Under Development, select **REST**.

Topics

- [About RESTful Web Services](#)
- [RESTful Services Terminology](#)
- [About the Overview Page](#)
- [About the AutoREST Page](#)
- [About the REST Search Toolbar](#)
- [Creating RESTful Web Services](#)
- [Securing RESTful Web Services](#)

4.1 About RESTful Web Services

Representational State Transfer (REST) is a style of software architecture for distributed hypermedia systems such as the World Wide Web. A service is described as RESTful when it conforms to the tenets of REST.

A RESTful service has the following characteristics:

- Data is modeled as a set of resources. Resources are identified by URIs.
- A small, uniform set of operations are used to manipulate resources (for example, GET, POST, PUT, DELETE).
- A resource can have multiple representations (for example, a blog might have an HTML representation and a RSS representation).
- Services are stateless and because it is likely that the client will want to access related resources, these should be identified in the representation returned, typically by providing hypertext links.

4.2 RESTful Services Terminology

The following are some major terms related to RESTful services.

- **RESTful service:** An HTTP web service that conforms to the tenets of the RESTful architectural style.

- **Resource module:** An organizational unit that is used to group related resource templates.
- **Resource template:** An individual RESTful service that is able to service requests for some set of URIs (Universal Resource Identifiers). The set of URIs is defined by the URI pattern of the Resource Template
- **URI pattern:** A pattern for the resource template. Can be either a route pattern or a URI template, although you are encouraged to use route patterns.
 - **Route pattern:** A pattern that focuses on decomposing the path portion of a URI into its component parts. For example, a pattern of `/objects/:object/:id?` will match `/objects/emp/101` (matches a request for the item in the emp resource with id of 101) and will also match `/objects/emp/` (matches a request for the emp resource, because the `:id` parameter is annotated with the `?` modifier, which indicates that the `id` parameter is optional).
 - **URI template:** A simple grammar that defines the specific patterns of URIs that a given resource template can handle. For example, the pattern `employees/{id}` will match any URI whose path begins with `employees/`, such as `employees/2560`.
- **Resource handler:** Provides the logic required to service a specific HTTP method for a specific resource template. For example, the logic of the GET HTTP method for the preceding resource template might be:

```
select empno, ename, dept from emp where empno = :id
```
- **HTTP operation:** HTTP (HyperText Transport Protocol) defines the standard methods that are supported in Oracle REST Data Services for building RESTful services: GET (retrieve the resource contents), POST (store a new resource), PUT (update an existing resource), and DELETE (remove a resource).

See Also

- [Structure of RESTful Web Services Video](#)
- [Getting Started with RESTful Services](#)

4.3 About the Overview Page

The Overview page provides an overview of the activity in the REST pages.

To navigate to the Overview page, do either of the following:

- In the Launchpad page, select the **Development** tab and click **REST**.
- Click **Selector**  to display the navigation menu. Under Development, select **REST**.

You can navigate to pages within REST by using the menu in the header: Overview, Modules, AutoREST and Security.

Use **Export Schema** to copy or download all the roles, privileges, modules, REST-enabled objects, and OAuth clients defined in the schema.

The Overview page consists of three sections.

Objects

Displays the total number of resource modules, schema objects with REST access, roles, privileges, OAuth clients, and pre-authenticated requests created. Click the respective card to navigate to the corresponding pages for modules, roles, privileges, OAuth clients, and pre-authenticated requests.

In the REST pages, the term "object" is used to refer to a module, template, handler, role, privilege, OAuth Client, or pre-authenticated request.

Security

Metadata Catalog: Displays whether access to the ORDS metadata catalog of the schema is protected. The metadata catalog lists all the services in the schema. The user authorization for the metadata catalog is provided at the time of enabling user access.

Modules: Displays the number of modules that are published out of the total modules created.

Module Security: Displays the number of modules that are protected by a privilege out of the total modules created. A partially secured privilege is where the module is protected both by a privilege and a pattern.

Recent Objects

Displays recently updated or created objects. Each entry has a context menu available at the end of the row. Click **Actions**  to access options such as Edit and Delete.

See Also

[Accessing Database Actions](#)

4.4 About the AutoREST Page

Use the Automatic Enabling of Schema Objects for REST Access (AutoREST) feature to enable REST access for database objects in the current schema.

Enabling REST access for a schema object such as table, view, function, procedure and package allows it to be accessed through RESTful services. See AutoREST in *Oracle REST Data Services Developer's Guide*

To navigate to the AutoREST page, in the REST Overview page, click **AutoREST** or select **AutoREST** from the menu in the header.

The AutoREST page consists of three parts:

- **Objects:** Displays the number of REST-enabled schema objects that are protected and require authorization out of the total number of schema objects enabled in the schema.
- **AutoREST:** Displays the number of tables and views, packages, procedures and functions that are REST enabled in the schema. Click a card to see the corresponding objects listed below.
- The bottom part of the page lists schema objects based on the selection made in AutoREST. You can view these objects in card view or grid view. The actions available in the context menu for an object are:

- **Edit:** Edit the Object Alias or Require Authentication fields for the REST enabled object. For a description of the fields, see [Enabling REST Access for a Database Object](#)
- **Disable:** Disable REST access for the object. See [Disabling REST Access for a Database Object](#)
- **Get Curl:** Generate cURL calls for a REST enabled object. See [Generating cURL Requests for a REST-Enabled Database Object](#)
- **Export OpenAPI:** Export the object as JSON Open API code.
- **OpenAPI View:** Display the object as a Swagger UI implementation. You can view and execute the handlers, pass parameters to the handlers, and copy or download the responses.

① See Also

- [About the REST Search Toolbar](#)
- [About the Overview Page](#)

4.4.1 Enabling REST Access for a Database Object

You can enable REST access for a table, view, materialized view, function, procedure or package.

To REST enable a schema object:

1. In the Navigator tab in the SQL page, right-click an object and select **REST**, then select **Enable**.
2. Enter the following fields:
 - **Object Alias:** Enter another name for the object. For security purposes, it is recommended to use a different name than the actual object name.
 - **Require Authentication:** Select this option to secure the REST endpoint.
 - **Authorization Role** and **Authorization Privilege:** These fields are automatically pre-filled.

Show code: Select this option to view the PL/SQL code equivalent of the Enable REST Access slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Enable** in the Enable REST Access slider.

3. Click **Enable**.

A REST-enabled object is indicated by a **REST Enabled** icon  in the Navigator tab.

4.4.2 Disabling REST Access for a Database Object

This section describes how to disable REST access for a database object.

1. In the SQL page, in the Navigator tab, right-click the database object that is REST enabled and select **Disable**.
2. You see a prompt to confirm. Click **OK**.

4.4.3 Generating cURL Requests for a REST-Enabled Database Object

For a REST-enabled database object, you can generate sample cURL calls for GET, POST, PUT and DELETE requests for the selected object. Use the Copy to Clipboard icon to copy the cURL code.

To view the cURL call for a specific method for the selected object:

1. In the SQL page, in the Navigator tab, right-click the REST-enabled object, select **REST** and then select **cURL command**.
2. Select the HTTP method from the left pane, enter values if required and click **Next** to see the cURL command.

For more information about each HTTP method, see [Examples in Oracle REST Data Services](#).

4.5 About the REST Search Toolbar

This section describes the various options available on the search toolbar in the REST pages. The options selected are saved for your session.



Search: Performs a search based on the text entered in the field. This field is not case sensitive. When you start typing in the Search field, the  icon appears. If you click the icon, the text that you typed is deleted and you can re-enter text again.

Filter by: Performs a search based on the filters selected. In the Filter by drop down, first select the column to apply the filter on, and then select the values to search for in that column.

Sort By: Sorts the data in ascending or descending based on the columns selected. The columns can be turned on or off. Reset restores the columns to the original state. The sort criteria selected is shown in the upper right corner as Sort By.

Page Size. First Page, Previous page, Next Page: Only displayed in the Card View. Enables you to select page size, and move between pages.

 **Reload:** Refreshes the information on the page.

There are two views available:

 **Open Card View:** This is the default display view. Information is displayed for each object in a card format. Each card displays associated attributes for the object along with the context menu, which list the actions available for the object.

 **Open Grid View:** Displays the information for each object in a single row of a table. The last column in each row contains the context menu icon, which shows the actions available for the object.

4.6 Creating RESTful Web Services

Create RESTful web services using the Modules, Templates and Handlers pages.

To create a RESTful web service, you need to:

- Define a resource module
- Define a resource template
- Define one or more resource handlers such as GET, PUT, POST or DELETE. Optionally, define parameters that you need to pass to the resource handler.

The following sections provide information on how to create resource modules, resource templates and resource handlers.

- [Managing Resource Modules](#)
- [Managing Resource Templates](#)
- [Managing Resource Handlers](#)
- [Example: Inserting a Record using a POST Handler](#)
- [Viewing Resource Handler Details and Managing Parameters](#)
- [Using Pre-Authenticated Requests to Access ORDS RESTful Services](#)

See Also

- [REST Workshop Overview Video](#)
- Creating a RESTful Service in *Oracle REST Data Services Quick Start Guide*


4.6.1 Managing Resource Modules

You can define, edit and delete resource modules in the Modules page.

To navigate to the Modules page, in the REST Overview page, click **Modules** in Objects, or from the menu in the header, select **Modules**.

The module attributes displayed by default in card view are shown in the following figure.

Module Published Status



The module status can be published or unpublished. If the status is unpublished, the icon is displayed in a red color.

Click the module name in card view to go to the Templates page for that module. You can also navigate to the Templates page through the context menu.

The actions available in the context menu are:

- [Creating a Resource Module](#)
- [Importing a Resource Module](#)
- [Editing a Resource Module](#)
- [Deleting a Resource Module](#)
- [Publishing/Unpublishing a Resource Module](#)
- [Exporting a Resource Module](#)
- [Viewing the Module in OpenAPI View](#)
- Navigating to the Templates page for a module

① See Also

[About the REST Search Toolbar](#)

4.6.1.1 Creating a Resource Module

This section describes how to create a resource module.

1. In the Modules page, click **Create Module**.
2. Enter the following fields. The fields with an asterisk (*) are mandatory:

Module Definition tab

- **Module Name:** Enter the name of the module. This field is case sensitive.
- **Base Path:** Enter the base of the URI for accessing the RESTful service. For example: hr/ means that all URIs starting with hr/ will be serviced by this resource module. Note the change in the **Preview URL** as you enter the base path.
- **Is Published:** Enable to make the RESTful service available for use.
- **Pagination Size:** Enter or select the number of results to return on each page based on a database query. The default value is 25.
- **Protected By Privilege:** Select a privilege to protect the module from the drop-down list.
- **Comments:** Enter descriptive comments.
- **Go to Module after creation:** Select this option to go to the Templates page after the module is created.

Origins Allowed tab

- Enter origins that are allowed to access the resource templates. Click **Add (+)** to add each origin. For example:

http://example1.org

https://*.example2.com

Show code: Select this option to view the PL/SQL code equivalent of the Create Module slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Module slider.

3. Click **Create**.

The new module is displayed on the Modules page. If you cannot see it, locate the module using the Search field.

4.6.1.2 Importing a Resource Module

This section describes how to import a resource module.

Note

You can import an OpenAPI version 2.0 or 3.0 document in JSON format only.

1. In the Modules page, at the top right, click **Import Module**.

2. Browse and select the JSON file to import. Click **Open**.

The Import Module slider is displayed with **Base Path** and **Preview URL** fields prefilled. Edit the Base Path as needed. Editing the base path field will also update the base path in the JSON format that is displayed in the slider.

3. Click **Import**.

When the module is imported, a privilege configuration is created and set to protect the module. The name of the privilege is generated by using the base path and adding "_priv" at the end of the base path string and also replacing every slash "/" in the base path to underscore "_".

4. Your module is ready for editing.

After your module is imported, the templates and handlers are created, but the module is not published. You need to supply SQL and PL/SQL blocks for each handler, and publish your module when it is ready for testing or production.

4.6.1.3 Editing a Resource Module

This section describes how to edit a resource module.

1. In the Modules page, for the specific module, click  and select **Edit**.
2. Edit the required fields and click **Save**. For a description of the fields, see [Creating a Resource Module](#).

4.6.1.4 Deleting a Resource Module

This section describes how to delete a resource module.

1. In the Modules page, for the specific module, click  and select **Delete**.
A prompt appears asking you to confirm.
2. Click **OK** to delete.

4.6.1.5 Publishing/Unpublishing a Resource Module

This section describes how to make a resource module available or not.

1. In the Modules page, for the specific module, click  and then select **Publish** or **Unpublish**.
You see a prompt to confirm.
2. Click **OK**.

4.6.1.6 Exporting a Resource Module

You can export the PL/SQL source code for the API or Open API (Swagger) JSON code for a resource module or RESTful service. If the module has templates, handlers, parameters, roles, and so on, these objects are also displayed and exported.

1. In the Modules page, for a specific module, click  and select **Export Module**.
2. Select **PL/SQL** to copy or export the PL/SQL source for the API. Select **OpenAPI** to copy or export the Open API (Swagger) JSON code for the module.

Note

A module exported using Open API is always imported with an unpublished status.

3. Click **Copy to Clipboard** to copy the code or click **Download** to download to your local computer.

4.6.1.7 Viewing the Module in OpenAPI View

This section describes how to open the module in Open API view.

The Open API view option is available for modules that are published and not protected.

1. In the Modules page, for the specific module, click  and select **OpenAPI View**.
2. The module is displayed as a Swagger UI implementation.

You can view and execute the handlers, pass parameters to the handlers, and copy or download the responses.

To learn more about Swagger UI implementation, refer to the official [Swagger UI](#) website.


4.6.2 Managing Resource Templates

You can define, edit and delete resource templates for a module in the Templates page.

To navigate to this page, in the Modules page, click  for a module and select **Templates**, or click the name of the module in card view.

At the top of the page, the module card is available with context menu options such as Edit, Delete, Publish and Export.

The templates attributes displayed by default in card view are shown in the following figure.



Click the template name in card view to go to the Handlers page. You can also navigate to the Handlers page through the context menu.

The actions available in the context menu are:

- [Creating a Resource Template](#)
- [Editing a Resource Template](#)
- [Deleting a Resource Template](#)
- Navigating to the Handlers page for the template

4.6.2.1 Creating a Resource Template

This section describes how to create a resource template.

1. In the Templates page for a module, click **Create Template**.
2. Enter the following fields. The fields with an asterisk (*) indicate that they are mandatory:

- **URI Template:** Enter the URI pattern for the resource template.

For example, a pattern of `/objects/:object/:id?` will match `/objects/emp/101` (matches a request for the item in the emp resource with id of 101) and will also match `/objects/emp/` (matches a request for the emp resource, because the `:id` parameter is annotated with the `?` modifier, which indicates that the id parameter is optional).

Note that the **Preview URL** changes as you type the URI template.

- **Priority:** Enter or choose the priority for how the resource template should be evaluated (1 is low priority and 9 is high priority).
- **HTTP Entity Tag Type:** Select the type of entity tag to be used by the resource template. An entity tag is an HTTP Header that acts as a version identifier for a resource.

Options include:

- **Secure HASH (default):** The contents of the returned resource representation are hashed using a secure digest function to provide a unique fingerprint for a given resource version.
- **Query:** Define a query manually that uniquely identifies a resource version. A manually defined query can often generate an entity tag more efficiently than hashing the entire resource representation.
- **None:** Do not generate an entity tag.
- **Comments:** Enter descriptive comments if any.
- **Go to Handlers after creation:** Select this option to go to the Handlers page after the template is created.

Show code: Select this option to view the PL/SQL code equivalent of the Create Template slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Template slider.

3. Click **Create**.

The new template for the module is added to the Templates page.

If "Go to Handlers after creation" is selected, you can use the module name in the breadcrumbs at the top of the page to see the list of templates for the module.

4.6.2.2 Editing a Resource Template

This section describes how to edit a resource template.

1. In the Templates page, for a specific template, click  and then select **Edit**.
2. Edit the required fields and click **Save**.

For a description of the fields, see [Creating a Resource Template](#).

4.6.2.3 Deleting a Resource Template

This section describes how to delete a resource template.

1. In the Templates page, for a specific template, click  and then select **Delete**.
You are prompted to confirm.
2. Click **OK** to delete.


4.6.3 Managing Resource Handlers

You can create, edit and delete resource handlers for a template in the Handlers page.

To navigate to this page, in the Templates page, click  for a template and select **Handlers**, or click the name of the template in card view.

At the top of the page, the template card is available with context menu options such as Edit and Delete.

The handler attributes displayed by default in card view are shown in the following figure.



The HTTP Method can be one of the following: GET, PUT, POST and DELETE.

You can navigate to the HTTP method page using the context menu. There are three tabs displayed in the HTTP method page: Source, Parameters and Pre-Authenticated Requests.

To create a resource handler, see [Creating a Resource Handler](#).

Note

This option is not available if the template has all four handlers created (GET, POST, PUT, DELETE).

The actions available in the context menu are:

- **Edit:** See [Editing a Resource Handler](#)

- **Delete:** See [Deleting a Resource Handler](#)
- **Details:** See [Viewing Resource Handler Details and Managing Parameters](#)
- **Pre-Authenticated Requests:** See [Pre-authenticated Requests](#)

4.6.3.1 Creating a Resource Handler

This section describes how to create a resource handler.

1. In the Handlers page, click **Create Handler**.
2. Specify the properties of the resource handler. The specific options depend on the method type.

Handler Definition Tab

- **Method:** Enter the HTTP request method for this handler: GET (retrieves a representation of a resource), POST (creates a new resource or adds a resource to a collection), PUT (updates an existing resource), or DELETE (deletes an existing resource). Only one handler for each HTTP method is permitted.
- **Items Per Page:** Enter or choose the default pagination size, or the number of rows to return for a database query. This option is only available for GET handlers. If this value is not defined, the number of items per page is the one defined in the module.
- **Source Type:** Select the source implementation for the selected HTTP method. The default is collection query.
 - **Collection Query:** Executes a SQL query and transforms the result set into a **JSON** representation. Available when the HTTP method is GET.
 - **Collection Item:** Executes a SQL query returning one row of data into a **JSON** representation. Available when the HTTP method is GET.
 - **Media:** Executes a SQL Query conforming to a specific format and turns the result set into a **binary** representation with an accompanying HTTP Content-Type header identifying the Internet media type of the representation.
 - **PL/SQL:** Executes an anonymous PL/SQL block and transforms any OUT or IN/OUT parameters into a **JSON** representation.
- **Source:** Enter the SQL query or the PL/SQL block for the HTTP method. In the handler details page, you can test the SQL or PL/SQL handler code in the handler editor.
- **Go to Handler after creation:** Select this option to go to the details page after the handler is created.

MIMEs Allowed tab

Identifies the type of information included in the HTTP Request Body. For example, if the POST Handler is expecting JSON in the request, add "Application/JSON" as the MIME type.

Show code: Select this option to view the PL/SQL code equivalent of the Create Handler slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Handler slider.

3. Click **Create**.

The handler is displayed on the Handlers page for the specific template.

You can test the REST service endpoint using a REST client or a command-line tool such as cURL.

① See Also

About cURL and Testing RESTful Services in *Oracle REST Data Services Developer's Guide*.

4.6.3.2 Editing a Resource Handler

This section describes how to edit a resource handler.

1. In the Handlers page, for a specific handler, click  and then click **Details**.
2. Edit the required fields and click **Save**.

For a description of the fields, see [Creating a Resource Handler](#).

4.6.3.3 Deleting a Resource Handler

This section describes how to delete a resource handler.

1. In the Handlers page, for the specific handler, click  and then click **Delete**.
You are prompted to confirm.
2. Click **OK** to delete.

4.6.3.4 Viewing Resource Handler Details and Managing Parameters

You can view resource handler details and manage parameters in the *HTTP method* page.

To navigate to this page, for a specific handler, click  and then select **Details** or click the name of the handler in card view.

At the top of the page, the handler card is available with context menu options such as Edit and Delete.

The Handler *HTTP method* page has two major sections:

- A code editor for executing the SQL or PL/SQL source code that was defined in the resource handler. You can execute the code, save, undo, redo, download and clear output. For more information about the editor, see [Executing SQL Statements in the Code Editor](#).

Implicit parameters used in REST service handlers are displayed in a scrolling list towards the right side of the editor. See [Implicit Parameters](#)

- A section for managing parameters required for running the HTTP method.

By default, the parameters are displayed in grid view. The attributes displayed on a parameter in card view are shown in the following figure.



The actions available in the context menu are:

- [Creating a Parameter](#)
- [Editing a Parameter](#)
- [Deleting a Parameter](#)

① **See Also**

[About the REST Search Toolbar](#)

4.6.3.4.1 Creating a Parameter

This section describes how to create a parameter for a resource handler.

1. In the handler *HTTP method* page, click **Create Parameter**.
2. Enter the following fields. The fields with an asterisk (*) are mandatory:
 - **Parameter Name:** Enter the name of the parameter, as it is named in the URI Template or HTTP Header. The name defines how the parameter is identified in the incoming request or how the parameter is named in the Response.
 - **Bind Variable Name:** Enter the name of the parameter, as it will be referred to in SQL. If NULL is specified, then the parameter is unbound.
 - **Source Type:** Select the type that is identified if the parameter originates in the URI Template or HTTP Header.
 - **Parameter Type:** Select the native type of the parameter.
 - **Access Method:** Select the parameter access method. Indicates if the parameter is an input value, output value, or both.

Show code: Select this option to view the PL/SQL code equivalent of the Create Parameter slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Parameter slider.

3. Click **Create**.

The parameter is added in the Parameters section.

4.6.3.4.2 Editing a Parameter

This section describes how to edit a parameter.

1. In the Parameters section, for the specific parameter, click  and then select **Edit**.
2. Edit the required fields and click **Save**.

For a description of the fields, see [Creating a Parameter](#).

4.6.3.4.3 Deleting a Parameter

This section describes how to delete a parameter.

1. In the Parameters section, for the specific parameter, click  and then select **Delete**.
You are prompted to confirm.
2. Click **Yes** to delete.

4.6.3.4.4 Implicit Parameters

This section describes the implicit parameters in the Resource Handler Details page. These parameters are automatically added to the resource handlers.

Table 4-1 Implicit Parameters

Name	Type	Access Mode	HTTP Header	Description
:body	BLOB	IN	N/A	Specifies the body of the request as a temporary BLOB.
:body_text	CLOB	IN	N/A	Specifies the body of the request as a temporary CLOB.
:content_type	VARCHAR	IN	Content-Type	Specifies the MIME type of the request body, as indicated by the Content-Type request header.
:current_user	VARCHAR	IN	N/A	Specifies the authenticated user for the request. If no user is authenticated, then the value is set to null.

Table 4-1 (Cont.) Implicit Parameters

Name	Type	Access Mode	HTTP Header	Description
:forward_location	VARCHAR	OUT	X-ORDS-FORWARD-LOCATION	Specifies the location where Oracle REST Data Services must forward a GET request to produce the response for this request.
:fetch_offset	NUMBER	IN	N/A	Specifies the zero-based offset of the first row to be displayed on a page.
:fetch_size	NUMBER	IN	N/A	Specifies the maximum number of rows to be retrieved on a page.

Table 4-1 (Cont.) Implicit Parameters

Name	Type	Access Mode	HTTP Header	Description
:page_offset	NUMBER	IN	N/A	Specifies the zero based page offset in a paginated request. Note: The :page_offset parameter is deprecated. Use :row_offset parameter instead.
:page_size	NUMBER	IN	N/A	Specifies the maximum number of rows to be retrieved on a page. The :page_size parameter is deprecated. Use :fetch_size parameter instead.
:row_offset	NUMBER	IN	N/A	Specifies the one-based index of the first row to be displayed in a paginated request.

Table 4-1 (Cont.) Implicit Parameters

Name	Type	Access Mode	HTTP Header	Description
:row_count	NUMBER	IN	N/A	Specifies the one-based index of the last row to be displayed in a paginated request.
:status_code	NUMBER	OUT	X-ORDS-STATUS-CODE	Specifies the HTTP status code for the request.

For more information about the implicit parameters, see *Oracle REST Data Services Developer's Guide*.

4.6.3.5 Using Pre-Authenticated Requests to Access ORDS RESTful Services

To access protected RESTful services, you can create and revoke pre-authenticated requests.

Pre-authenticated requests enable you to access protected ORDS RESTful services without user credentials. When you create pre-authenticated request, a unique URL is generated. You can provide this URL to an external user to interact with the particular RESTful entity using standard HTTP tools.

For more information about Pre-Authenticated Requests, see Oracle REST Data Services Pre-Authenticated Requests in the *Oracle REST Data Services Developer's Guide*.

To navigate to the Pre-Authenticated Requests page, you can do either of the following:

- From the Overview page, click the **Pre-Authenticated Requests** tile.
- From the menu bar, expand **Security** and select **Pre-Authenticated Requests**.

Topics:

- [Creating a Pre-Authenticated Request](#)
- [Viewing a Pre-Authenticated Request](#)

4.6.3.5.1 Creating a Pre-Authenticated Request

To create a pre-authenticated request for an existing handler:

1. From the context menu for an existing handler, select **Create New Pre-authenticated Request**.

The Create Pre-authenticated Request panel appears.

2. In the **Token Expiration** field, enter the token expiry duration in seconds.

Select **Show Code** to view the underlying source code for creating the pre-authenticated request.

Click **Create**.

The token, alias, and relative link to access the resource are displayed. You must keep track of the token and alias as you cannot obtain their values later. If you forget the values, you can revoke the token and create a new one again.

4.6.3.5.2 Viewing a Pre-Authenticated Request

From the context menu for an existing handler, select **Pre-Authenticated Requests** and then select **Details**.

The pre-authenticated requests for the specific handler appears.

4.6.3.5.3 Revoking a Pre-Authenticated Request

To revoke the pre-authenticated request for a handler:

1. From the context menu for an existing handler, select **Pre-Authenticated Requests** and then select **Details**.

The pre-authenticated requests for the specific handler appears.

2. From the context menu for the pre-authenticated request, select **Revoke**.

You are prompted to confirm.

3. Click **Confirm** to revoke the request.

4.6.4 Example: Inserting a Record using a POST Handler

The following example illustrates how to insert a record in the DEPT table.

Prerequisites

- Using the worksheet in the SQL page, create a **DEPT** table.

```
CREATE TABLE DEPT(  
  DEPTNO    number(2,0),  
  DNAME     varchar2(14),  
  LOC       varchar2(13),  
  CONSTRAINT PK_DEPT PRIMARY KEY (DEPTNO)  
)
```

- Create a module named **example**. See [Creating a Resource Module](#)
- Create a template named **emp/** for the module. See [Creating a Resource Template](#)

To insert a record:

1. In the Handlers page, click **Create Handler**.
2. Enter the following details:
 - In the **Method** field, select **POST**.
 - In the **Source Type** field, select **PL/SQL**.

- In the **Source** field, enter the following PL/SQL block:

```
DECLARE
  id number;
BEGIN
  INSERT INTO DEPT (DEPTNO, DNAME, LOC) VALUES
  (:DEPTNO, :DEPTNAME, :DEPTLOC) RETURNING DEPTNO INTO id;
  :status_code := 201;
  :forward_location := id;
END;
```

where

- RETURNING DEPTNO INTO id returns the value assigned to DEPTNO into the variable id.
- :status_code is an implicit parameter that is assigned 201 to indicate that the resource is created.
- :forward_location is an implicit parameter that is assigned the id that contains the value of DEPTNO. This parameter specifies the location to forward a GET request to produce the response for the POST request.
- Ensure **Go to Handler page after creation** is selected.
- In the **MIMEs Allowed** Tab, select **application/json** from the drop-down list and click **Add +**.
- Click **Create**.

The Handler POST page appears.

You can test the source code by clicking **Run Statement**  under the Source tab.

- Create a GET resource handler to produce a response for the POST request.

In the Example module, create a template named **emp/:id**.

Create a GET resource handler where **Source Type** is **Collection Item** and **Source** contains the following query:

```
SELECT * FROM DEPT WHERE DEPTNO = :id
```

- Test the RESTful service endpoint using the following command in cURL. You can copy the URL (<http://xyz.us.comp.com:1234/ords/pdbdba/example/emp>) using the Copy to Clipboard icon in the POST handler card at the top.

```
curl --location --request POST --header "Content-Type: application/json"
--data '{"DEPTNO": 54, "DEPTNAME": "HR", "DEPTLOC": "America" }' 'http://
xyz.us.comp.com:1234/ords/pdbdba/example/emp/'
```

The output is:

```
{"deptno":54,"dname":"HR","loc":"America","links":[{"rel":"collection",
"href":"xyz.us.comp.com:1234/ords/pdbdba/example/emp/"}]}
```

5. In the SQL page, check the DEPT table to see if the new record has been inserted by using the following statement:

```
SELECT * FROM DEPT;
```

4.7 Securing RESTful Web Services

Define roles, privileges and OAuth Clients to ensure authentication and authorization are required for accessing RESTful web services.

To protect a RESTful web service, you need to:

- Create a role
- Create a privilege selecting the role and modules or resources to protect

To enable access to a protected RESTful service using the OAUTH2 Workflow, create an OAuth client using the role and privilege created for protecting the RESTful service.

The following sections provide information on how to create roles, privileges and OAuth clients:

- [Managing Roles](#)
- [Managing Privileges](#)
- [Managing OAuth Clients](#)
- [Examples](#)

See Also

- [REST Workshop Overview Video](#)
- Tutorial for Protecting and Accessing Resources

4.7.1 Managing Roles

You can create, edit and delete roles for RESTful services in the Roles page.

To navigate to the Roles page, from the REST Overview page, click **Roles** in Objects, or from the menu in the header, select **Security** and then select **Roles**.

The actions available in the context menu are:

- [Creating a Role](#)
- [Editing a Role](#)
- [Deleting a Role](#)
- [Viewing Assigned Privileges](#)

See Also

About Users and Roles in *Oracle REST Data Services Developer's Guide*.

4.7.1.1 Creating a Role

Create a role with a specific name. After the role is created, you can associate it with a privilege.

1. In the Roles page, click **Create Role**.
2. In the Role Name field, enter the name of the role to create.

Show code: Select this option to view the PL/SQL code equivalent of the Create Role slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Role slider.

3. Click **Create**.

The new assigned role is displayed on the Roles page.

4.7.1.2 Editing a Role

This section describes how to edit a role.

1. In the Roles page, for the specific role, click **Actions**  and select **Edit**.
2. Enter the changes required and click **Save**.

For a description of the fields, see [Creating a Role](#).

4.7.1.3 Deleting a Role

This section describes how to delete a role.

1. In the Roles page, for the specific role, click **Actions**  and select **Delete**.
You are prompted to confirm.
2. Click **Yes** to delete.

4.7.1.4 Viewing Assigned Privileges

This section describes how to view the privileges associated with a role.

In the Roles page, for the specific role, click **Actions**  and select **Details**. The privileges assigned to the role are displayed.

4.7.2 Managing Privileges

You can create, edit and delete privileges for RESTful services in the Privileges page.

A privilege defines the set of roles, at least one of which an authenticated user must possess to access a RESTful service protected by a privilege.

To navigate to the Privileges page, from the REST Overview page, click **Privileges** in Objects, or from the menu in the header, select **Security** and then select **Privileges**.

The privilege attributes displayed by default in card view are shown in the following figure.

The screenshot shows a web interface for configuring a privilege. The form contains the following elements:

- oracle.demo.example** → Name
- Demo** → Label
- Privilege example** → Description
- HR** (selected role)
- Roles: 1** (Number of Roles Granted)
- Modules: 1** (Protected Modules)
- Patterns: 0** (Protected Patterns)
- Context Menu** (indicated by a red arrow pointing to a three-dot menu icon)

The actions available in the context menu are:

- [Creating a Privilege](#)
- [Editing a Privilege](#)
- [Deleting a Privilege](#)

4.7.2.1 Creating a Privilege

This section describes how to create a privilege.

1. In the Privileges page, click **Create Privilege**.
2. Enter the following fields. The fields with an asterisk (*) are mandatory:
 - **Label:** Enter the name of the privilege in a way that it is easy to understand.
 - **Name:** Enter a unique name for the privilege.
 - **Description:** Enter a brief description of the purpose of the privilege.
 - **Comments:** Enter comments.
 - **Roles:** Enter one or more roles assigned to the privilege.
 - **Protected Modules:** Select the modules to protect.
 - **Protected Resources:** Instead of Protected Modules, use the Protect Resources tab to apply security based on a URI pattern.

Show code: Select this option to view the PL/SQL code equivalent of the Create Privilege slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Privilege slider.

3. Click **Create**.

The new privilege is displayed on the Privileges page.

4.7.2.2 Editing a Privilege

This section describes how to edit a privilege.

1. In the Privileges page, for the specific privilege, click **Actions** and select **Edit**.
2. Make the required changes and click **Save**.

For a description of the fields, see [Creating a Privilege](#).

4.7.2.3 Deleting a Privilege

This section describes how to delete a privilege.

1. In the Privileges page, for the specific privilege, click **Actions** and select **Delete**.
2. You are prompted to confirm. Click **Yes**.

4.7.3 Managing OAuth Clients


Using OAuth 2.0-based authentication, you can ensure that your RESTful web services are accessed only by specific users or clients.

OAuth 2.0 is a standard Internet protocol that defines flows to provide conditional and limited access to a RESTful API. For more information, see [OAuth-Based Authentication](#).

You can create, edit and delete OAuth Clients in the OAuth Clients page.

To navigate to the OAuth Clients page, from the REST Overview page, click **Clients** in Objects, or from the menu in the header, select **Security** and then select **OAuth Client**.

The OAuth Client attributes displayed by default in card view are shown in the following figure.



To create an OAuth client, see [Creating an OAuth Client](#).

The actions available in the context menu are:

- **Edit:** See [Editing an OAuth Client](#)
- **Export:** See [Exporting an OAuth Client](#)
- **Delete:** See [Deleting an OAuth Client](#)
- **Get Bearer Token:** This option is displayed only if the client is created with a non-encrypted client secret. It provides the access token to call the RESTful service for Client Credentials and Implicit OAuth grant types. See [Creating an OAuth Client Using the Client Credentials Grant Type](#)
- **Manage Secrets:** Applicable for Client Credentials and Auth Code grant types. See [Managing Secrets](#)
 - **Rotate:** Removes the existing client secret and generates a new one for the OAuth client.
 - **Revoke:** Removes the client secret associated with the OAuth client.
- **Auth Details:** Displays the Unique Value and Authorization URI for the Auth Code OAuth grant type. See [Creating an OAuth Client Using the Auth Code Grant Type](#)

 See Also

Tutorial for Protecting and Accessing Resources

4.7.3.1 Creating an OAuth Client

Creates the OAuth Client and grants the required roles and privileges.

1. In the OAuth Clients page, select **Create OAuth Client**.
2. Enter the following fields. The fields with an asterisk (*) are mandatory:

Client Definition tab

- **Grant type:** Select the authorization grant type. The options are Client_Cred, Auth Code, or Implicit.
For more information about these grant types, see OAuth Flows in *Oracle REST Data Services Developer's Guide*.
- **Name:** Name of the client.
- **Description:** Description of the purpose of the client.
- **Redirect URI:** Enter the client-controlled URI to which redirect containing an OAuth access token or error will be sent.
- **Support URI:** Enter the URI where end users can contact the client for support.
Example: [http:// www.myclientdomain.com/support/](http://www.myclientdomain.com/support/)
- **Support Email:** Enter the email where end users can contact the client for support.
- **Logo:** Upload your logo in JPG, BMP or SVG file format. Ensure that the logo is less than 100 KB in size.
- **Roles:** Select roles to be granted.
- **Allowed Origins:** Add the list of URL prefixes.
- **Privileges:** Select privileges that the client wants to access.

Show code: Select this option to view the PL/SQL code equivalent of the Create OAuth Client panel. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create OAuth Client panel.

3. Click **Create**.

The OAuth Client registered is displayed on the OAuth Clients page.

A preview of the client secret is displayed. Secrets are applicable only for Client Credentials and OAuth Code grant types. Copy the secret value as this is the only instance when it is shown. If you forget the client secret, you can use the Rotate Secret action to change it. See [Managing Secrets](#).

The Client ID value represents the secret credential for the OAuth client. Click **Show/Hide**  to see the values.

Test the secured REST service endpoint using a REST client or the cURL command line tool.

① See Also

Tutorial for Protecting and Accessing Resources in *Oracle REST Data Services Developer's Guide*

4.7.3.2 Editing an OAuth Client

This section describes how to edit an OAuth Client.

1. In the OAuth Clients page, for the specific client, click **Actions** ⋮ and select **Edit**.
2. Edit the required fields and click **Save**.

For a description of the fields, see [Creating an OAuth Client](#).

4.7.3.3 Deleting an OAuth Client

This section describes how to delete an OAuth Client.

1. In the OAuth Clients page, for the specific client, click **Actions** ⋮ and select **Delete**.
2. You are prompted to confirm. Click **Yes**.

4.7.3.4 Exporting an OAuth Client

This section describes how to export an OAuth Client.

1. In the OAuth Clients page, for the specific client, click **Actions** ⋮ and select **Export**.
2. In the OAuth Client panel, click the **Copy** icon or **Download** to copy or download the OAuth Client information.

4.7.3.5 Managing Secrets

After a client secret is generated for an OAuth client, you can rotate or revoke the secret when required. The Rotate and Revoke options are available in the specific OAuth client's context menu.

① Note

With the deprecation of OAUTH and OAUTH_ADMIN PL/SQL packages (see [ORDS_SECURITY PL/SQL Package Reference](#)), the OAuth client secret creation process changes for Client Credentials and Auth Code grant types.

For OAuth clients that have been created with a non-encrypted secret, the label Legacy is displayed on the card.

Rotate a Client Secret

Removes the existing secret and creates a new one. This is useful when you cannot remember the existing client secret value.

- If a client secret exists for the OAuth client, click **Rotate** to remove the existing one and generate a new client secret.

- If a client secret is revoked and there is no secret value assigned, click **Register** to generate a client secret for the OAuth client.

Revoke a Client Secret

Removes the existing client secret.

Select **Revoke sessions** to remove all existing client sessions.

4.7.4 Examples

This section provides a few examples on creating an OAuth client with different grant types.

Topics:

- [Creating an OAuth Client Using the Client Credentials Grant Type](#)
- [Creating an OAuth Client Using the Auth Code Grant Type](#)

4.7.4.1 Creating an OAuth Client Using the Client Credentials Grant Type

This section describes how to create an OAuth Client using the Client Credential grant type.

Create an OAuth Client for the created module "example" in [Example: Inserting a Record using a POST Handler](#). The endpoint for the RESTful service is `http://xyz.us.comp.com:1234/ords/pdbdba/example/emp/`.

Prerequisites

Create a role named HR Admin. See [Creating a Role](#)

Create a privilege named Example.HR. See [Creating a Privilege](#)

1. In the OAuth Clients page, click **Create OAuth Client**.
2. Enter the following fields:
 - In the Grant type field, select **CLIENT_CRED**.
 - Enter Name, Description, Redirect URI, Support URI and Support Email for your OAuth Client.
 - In the Roles tab, add the created role (HR Admin).
 - In the Privileges tab, add the created privilege (Example.HR).
 - Click **Create**.


The new OAuth Client card appears on the OAuth Clients page.

3. Using cURL, test the service endpoint without OAuth credentials.


```
curl --location --request POST --header "Content-Type: application/json"
--data '{"DEPTNO": 55, "DEPTNAME": "Sales", "DEPTLOC": "Australia" }'
'http://xyz.us.comp.com:1234/ords/pdbdba/example/emp/'
```

You see an error.

4. Test the endpoint with OAuth credentials:
 - a. To get an access token, select **Get Bearer Token** from the context menu in the OAuth Client card.



The OAuth Token dialog appears. Use **Copy to Clipboard**  to copy the token.



- b. In cURL, use the bearer access token in the following statement to request the resource:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer
AMccwInt99gm9kxDs_w1DA" --location --request POST --data
 '{"DEPTNO": 55, "DEPTNAME": "Sales", "DEPTLOC": "Australia"}' 'http://
xyz.us.comp.com:1234/ords/pdbdba/example/emp/'
```

The output shows:

```
{"deptno":55,"dname":"Sales","loc":"Australia","links":
 [{"rel":"collection",
 "href":"http://xyz.us.comp.com:1234/ords/pdbdba/example/emp/"}]}
```

5. In the SQL page, you can verify the new record that has been inserted in the DEPT table by using the following statement:

```
SELECT * FROM DEPT;
```

4.7.4.2 Creating an OAuth Client Using the Auth Code Grant Type

This section describes how to create an OAuth Client using the Auth Code grant type.


1. In the OAuth Clients page, click **Create OAuth Client**.
2. Enter the following fields:
 - In the **Grant type** field, select **Auth Code**.
 - Enter **Name**, **Description**, **Redirect URI**, **Support URI** and **Support Email** for your OAuth Client.
 - In the **Roles** tab, add the relevant roles.
 - In the **Privileges** tab, add the relevant privileges.
 - Click **Create**.

The new OAuth Client card appears on the OAuth Clients page.

3. In the OAuth Client card, click the **Actions** icon and select **Auth Details**.

The Unique Value and Authorization URI appears.

4. In the Authorization URI field, click the **Open in New Tab** icon. A new window is displayed with an Unauthorized error message along with a Sign In link.
5. Click **Sign in** and enter your login credentials again.
6. You are prompted to approve the request to access the protected URI. Click **Approve**.



5

The Liquibase Page

The Liquibase page displays information of all database deployments made in the current schema. A deployment consists of changesets, which is a list of sequential changes to be applied to the database.

To navigate to the Liquibase page, do either of the following:

- In the Launchpad page, select the **Development** tab and click **Liquibase**.
- Click **Selector**  to display the navigation menu. Under **Development**, select **Liquibase**.

If you have not made any deployments into the schema and you access the Liquibase page, no data appears.

Topics:

- [Generate a Deployment](#)
- [About the Liquibase User Interface](#)

5.1 Generate a Deployment

First, you have to make changes in the current schema such as creating a table, procedure or any database object.

Then, create a folder that will contain the changesets for the deployment. It is a good practice to create a different folder for each deployment. This ensures that you have information of all the deployments, in case you want to roll back to a previous version.

In the following example, Oracle SQLcl is used to deploy the changes.

1. Open the terminal and navigate to the path of the folder that you will use to save the changes made to the database (changelog files).
2. After you navigate to the correct path, log in to the schema you want to capture with the following command:

```
sqlcl <username>/<password>@<host>:<port>/<servicename>
```

3. You can do one of the following:
 - Deploy the entire schema with the following command:

```
lb genschema
```

This command creates a controller.xml file that includes all the changesets for the schema.

- Deploy a specific object using the following command:

```
lb genobject -type <object_type> -name <object_name>
```

This command generates the changelogs (which are XML files) that contain the changes made to the database in the current folder.

4. Log in to the schema where you want to add these changes (in this example, XYZ) with the following command:

```
sqlcl xyz/<password>@<host>:<port>/<servicename>
```


5. Run the following command:

```
lb update -changelog controller.xml (or the file name of your changelog)
```

6. Log in to Database Actions on the schema where you deployed the changes and the Liquibase page displays information about the deployment.

5.2 About the Liquibase User Interface

In the Liquibase page, the card at the top displays the Liquibase status. The deployments are displayed as cards below, as shown in the following figure.



Liquibase Status


Liquibase runs only one deployment at a time. When there is no deployment running, the top Liquibase card indicates that Liquibase is currently unlocked and displays the ID of the last deployment.

If there is a deployment running at the time, the card indicates that Liquibase is locked, and displays the name of the resource blocking it and the time at which the deployment started.



Deployments


The card for each deployment provides details such as the schema, deployment id, start and complete deployment date and time, and number of changesets run in the deployment. To open a specific deployment page, click the title of the card, or click **Open Deployment Changes** in the context menu.



Changesets

In a deployment page, the changesets that have been run for that specific deployment is displayed below. The Current Deployment card at the top shows the deployment to which the changesets belong.

Each changeset card displays the changeset name, the date and time when the changeset started deploying, and the schema where the change was deployed.



The context menu for a changeset consists of the following options:

- **Previous Object State XML:** Displays an XML showing the state of the object (table, procedure) before the changeset was applied. If the object did not exist previously on the database (example, a table that was created in the deployment), the following message is displayed: `xml code not available....`
- **Show Executed SQL:** Displays the SQL code that was executed when the changeset was deployed.

6

The Charts and Dashboards Page

Charts enable you to create charts from the database. The chart is constructed using the input SQL command. Dashboards enable you to group charts together to create reports.

To navigate to the Charts and Dashboards page, do either of the following:

- In the Launchpad page, select the **Development** tab and click **Charts**.
- Click **Selector**  to display the navigation menu. Under Development, select **Charts**.

Topics

- [Overview](#)
- [Creating or Editing a Chart](#)
- [Example: Creating and Editing a Chart](#)
- [Creating or Editing a Dashboard](#)

6.1 Overview

This section describes the layout of the main Charts and Dashboards page.

To navigate to the Charts and Dashboards page, do either of the following:

The upper right corner has two icons:

- **Tour:** Provides a guided tour of charts and dashboards highlighting salient features and providing information that is useful if you are new to the interface.
- **Create:** Creates a new chart or dashboard. See [Creating or Editing a Chart](#) and [Creating or Editing a Dashboard](#)

This page consists of two sections:

- **Charts and Dashboards:** Displays the total number of charts and dashboards created. Click the respective card to view the charts or dashboards created in the section below.
- Displays the charts or dashboards created in the default card format.

The actions available for a chart or dashboard are:

- **View Chart/Dashboard:** To view the chart or dashboard on a different page.
- **Edit:** To make changes to the chart or dashboard.
- **Unpublish:** To make the chart or dashboard unavailable for use.
- **Go to module definition:** To show the module definition in the REST Modules page.
- **Delete:** To delete the chart or dashboard.

6.2 Creating or Editing a Chart

You can create or edit a chart to visualize data using SQL commands.

To create a new chart, click **Create Chart**.

To edit a chart, click the context menu icon and select **Edit**.

Enter the following details:

1. **Name:** Enter a name for the chart.
2. **Description:** Enter a description for the chart.
3. **Protected by Privilege:** Select an available privilege from the drop-down list to enable only those with this privilege to securely access the chart.
4. **Published:** When enabled, you can share and access the chart externally.

Note

You can also protect a chart by protecting the corresponding module using the REST pages. For more information, see [Editing a Resource Module](#). When a chart is protected, a user must authenticate before viewing the chart.

5. Click **Next**.
6. Enter the SQL statement in the command editor on the left. Ensure that you do not end the SQL statement with a semicolon.

The chart settings are automatically updated in the **Chart Definition** section on the right.

7. Edit/Update the chart settings in the **Chart Definition** section. The fields in the mapping section varies according to the **Chart Type** selected.

When the **Chart Type** selected is one of the following: *Area Chart*, *Bar Chart*, *Line Chart*, or *Line with Area Chart*, the mappings displayed are as follows:

- **Orientation:** Choose between horizontal and vertical orientation types from the drop-down list.
- **X axis label and Y axis label:** Optionally enter labels for X axis and Y axis.
- **Label:** The label is auto-generated from the SQL statement which represents the X axis.
- **Value:** The Value is auto-generated from the SQL statement which represents the Y axis (the values to be plotted on the chart).
- **Refresh Columns:** Refresh the data of all the columns selected in the SQL statement.
- **Coordinate System:** Choose between Cartesian and Polar coordinate systems from the drop-down list.
- **Sorting:** Choose a sorting type (ascending, descending, off) for plotting the values. If selected, the Sorting value overrides the order specified by the ORDER BY clause in the SQL statement. For a large amount of data, Oracle recommends using the ORDER BY clause to sort data.

When the **Chart Type** selected is one of the following: *Pie Chart*, *Pyramid Chart*, or *Funnel Chart*, the mappings displayed are as follows:

- **Style:** Select from either 2D or 3D styles.

- **Series ID:** The Series ID is auto-generated from the SQL statement which represents the category to be plotted.
- **Value:** The Value is auto-generated from the SQL statement which represents the values to be plotted for each category.

Preview/Auto-preview: Click preview to view the changes made to the chart. If Auto-preview is checked, the changes to the chart are refreshed automatically.

8. The **Chart** tab at the bottom displays the chart.
9. The **Data** tab at the bottom displays the data resulting from the SQL command.
10. Click **Create** to create the chart, or click **Save** when editing the chart.

Note

Due to browser resource restrictions, the preview of charts is limited to 3000 rows from your SQL query results. However, once you create a chart, all the data is rendered in the standalone view.

6.3 Example: Creating and Editing a Chart

This section provides an example of creating and editing a chart.

Creating a Chart

The following example shows how to create a chart to understand the cumulative salary earned by each job category.


1. Navigate to the **Charts** page from Database Actions homepage.
2. Click **Create Chart**.
3. Enter the details of the chart as follows:
 - **Name:** Total Salary by Job Category
 - **Description:** The cumulative salary earned by each job category.
 - **Chart Definition:**
 - **SQL Statement:**

```
SELECT SUM(SALARY), JOB_ID
FROM EMPLOYEES
GROUP BY JOB_ID
```


- **Chart Type:** Area Chart
- **Orientation:** Vertical
- **X axis label:** Job Category
- **Y axis label:** Total Salary
- **Label:** JOB_ID (VARCHAR2)
- **Value:** SUM(SALARY) (NUMBER)
- **Coordinate system:** Cartesian

— **Sorting:** Off

- Click **Preview** in the Mapping section to preview the chart below.




- Change the **Chart Type** to *Bar Chart*, **Sorting** to *Descending* and click **Preview** to view the updated chart.



- Click **Create** to create the chart. A confirmation notification is displayed and the chart card is visible in the Charts page.


- Click the Selector  and select **View Chart** to view the chart in a new tab.



Editing a Chart

The following example edits the chart created in the previous section.

- Click the Selector  and select **Edit** to edit the chart.



The screenshot shows the 'Charts' page in Oracle APEX. The chart 'Total Salary by Job Category' is selected, and a context menu is open. The menu options are: View Chart, Edit, Unpublish, Go to module definition, and Delete. The chart is published and unprotected.

- In the Mapping section, change the **Chart Type** to *Pie Chart* and **Style** to *3D*.

- Click **Preview** to view the updated chart.

Edit - Total Salary by Job Category

General Details

Name * Total Salary by Job Category Published *

Description * The cumulative salary earned by each job category.

Chart Definition

```

1 SELECT SUM(SALARY), JOB_ID
2 FROM EMPLOYEES
3 GROUP BY JOB_ID

```


Chart type * Pie Chart Style 3D

Series ID Column * JOB_ID (VARCHAR2) Value Column * SUM(SALARY) (NUMBE... ☐ Auto-preview **Preview**

Chart Data

Save Cancel

- Click **Save**.
- Click the Selector and select **View Chart** to view the updated chart in a new tab.




6.4 Creating or Editing a Dashboard

This section describes the steps to create or edit a dashboard.

To create a dashboard, in the Charts and Dashboard page, at the top right, click **Create** and then select **New Dashboard**.

To edit a dashboard, click the Actions icon and select **Edit**.

1. In the Create a Dashboard pane, enter or edit the following fields:
 - **Name:** Enter a name for the dashboard.
 - **Published:** Select this option to make the dashboard available for use.
 - **Description:** Enter a description for the dashboard.
 - **Protected by Privilege:** Select an available privilege from the drop-down list to enable only those with this privilege to securely access the dashboard.
 - **Go to Dashboard details after creation:** Select this option to go to the main Charts and Dashboards page after creating the dashboard.
2. Click **Create** or click **Save**, if editing. The Layout page is displayed
3. In the Layout page, you can add charts and design the layout for the dashboard. Click **Add Chart**.
 - In the **Select Chart** pane, for the **Selection** tab, select the chart to add from the drop-down list.
You can select the **Custom** option and enter a URL for the chart from your schema or from a different schema. The URL must point to the REST Module root, for example, `http://xxxx.oracle.com:1234/ords/hr/sdw/charts/Chart01/`. Also, you need to have the associated privileges to view the chart.
 - A preview of the selected chart is displayed. You can change the refresh rate of the chart, display the chart in Chart or tabular format, or share the chart by copying the URL.
 - In the **Settings** tab, select **Chart** and **Data Grid** to enable both views for the chart. By default, the Chart view is displayed in Preview.
4. Click **Create**.
The chart is added to the Layout section.



5. Add more charts as needed to the dashboard. The maximum is 3 or 4 based on the number of columns that you select from the Columns drop-down list.
 - You can reorder the charts by dragging them horizontally or vertically to move them to the required position.

- Use the Edit icon on a chart to edit the chart if needed. You can also change the chart by selecting from the drop-down list.
 - Drag the double-sided arrow at the right corner to resize the chart.
6. Click **Save and Preview** to save the layout and display a preview of the dashboard.

The dashboard is displayed on a new page. For each chart, you can change or customize the refresh rate, display the chart in Chart or tabular format, or share the chart by copying the URL.

7

The JSON Page

Note

Available for Oracle Database release 19c and later releases and only if you signed in as a database user with the **SODA_APP** role.

Use the JSON page in Database Actions to view and manage JSON collections, search for collection items using Query-by-Example (QBE), create JSON search indexes and display data guide diagrams for collections.

To navigate to the JSON page, do either of the following:

- In the Launchpad page, select the **Development** tab and click **JSON**.
- Click **Selector**  to display the navigation menu. Under **Development**, select **JSON**.

Topics

- [About the JSON User Interface](#)
- [Managing JSON Collections](#)
- [About Querying Documents in a Collection](#)
- [Creating Indexes for JSON Collections](#)
- [Creating Relational Views of JSON Documents](#)
- [Viewing the JSON Data Guide Diagram for a Collection](#)

See Also


[SODA Overview in Oracle Database Introduction to Simple Oracle Document Access \(SODA\)](#)

7.1 About the JSON User Interface

The JSON user interface consists of the left pane for listing and searching saved collections and the right pane for viewing and managing documents in a collection.

Listing and Searching JSON Collections

The following figure shows the main items in the left pane of the JSON page.



Select the appropriate option in the drop-down list to display saved JSON collections or recently accessed JSON collections.

The search functionality is not case-sensitive, retrieves all matching entries, and does not require the use of wildcard characters.


Click the Create Collection icon to create a new collection. See [Creating a Collection](#)

Right-click a collection name to open the context menu. The available options are:

- **Indexes:** Enables you to view existing JSON indexes and create search, functional or spatial indexes.
- **Drop:** Removes the collection from the database completely.

Viewing the Contents of a Collection

When a specific collection is selected in the left pane, the JSON documents that belong to the collection are displayed in the lower part of the right pane. The top part of the right pane contains the JSON editor where you can run queries for filtering or sorting the documents.



The two icons in the right corner are:

- **Tour** : Starts the JSON tour, which provides information about the features available.
- **Go to SQL** : Navigates to the SQL page.

The icons in the toolbar are:

- **New JSON Document** : Adds a new document to the collection. See [About Adding or Editing a JSON Document](#)
- **Import Document(s)** : Imports one or more existing JSON files from your local computer into the collection.
- **Delete All Documents in the list** : Deletes all JSON documents in the collection that match the current QBE search string. If the current string is {}, then all the documents in the collection are deleted.
- **Collection Details** : Enables you to view collection properties, JSON data guide (if created) and related statistics if they are gathered, size of search index, and page for managing JSON indexes.
- **Diagram** : Displays the the JSON data guide as a diagram in a hierarchical format. See [Viewing the JSON Data Guide Diagram for a Collection](#)
- **New Collection View** : Creates relational views of documents in a collection.

- **Run Query** : Filters documents using the QBE condition entered in the JSON editor. See [About Querying Documents in a Collection](#)
- **View Collection SQL** : Displays the JSON collection in SQL format. You can execute the SQL code, download it or copy to clipboard.

Each JSON document has the following icons:

- **Edit Document** : Edits the JSON document. See [About Adding or Editing a JSON Document](#)
- **Clone Document** : Creates a clone of the document.
- **Copy Document** : Copies the document to the clipboard.
- **Export Document** : Downloads the document as a .JSON file.
- **Delete Document** : Deletes the document.

7.2 Managing JSON Collections

You can add, view or drop collections, or browse, add, edit and delete JSON documents in a collection.

- **Create a Collection:** See [Creating a Collection](#)
- **Add or Edit Documents in a Collection:** See [About Adding or Editing a JSON Document](#)
- **View Documents in a Collection:** Select the collection name in the left pane. The documents in the collection are displayed in the right pane.
- **View Collection Details:** Select the specific collection in the left pane, and then click **Collection Details**  in the right pane toolbar to view collection properties, JSON data guide (if created) and related statistics if gathered, the size of search index, and the page for managing JSON indexes.
- **Drop a Collection:** Right-click a collection name in the left pane to open the context menu. Select **Drop** to remove the collection from the database completely.

7.2.1 Creating a Collection

This section describes how to create a collection.

1. When you navigate to the JSON page:
 - If there are no collections created, you see the JSON home page. Click **Create Collection**.
 - If there are existing collections available, then in the left pane, click **Create Collection** .
2. Enter the collection name, which is case sensitive.
3. Click **Create** to create the collection.
Click **Cancel** or press the **Esc** key to cancel the collection.

① See Also

[SODA Collection Metadata Components in Oracle Database Introduction to Simple Oracle Document Access \(SODA\)](#)


7.2.2 About Adding or Editing a JSON Document

You can add and edit JSON documents using the JSON editor.

To add a document, click **New JSON Document**  in the right pane toolbar for a specific collection. In the JSON editor, you can copy and paste the JSON document or use the Import icon to import the JSON document.

To edit a document, click **Edit Document**  in the document card.

The following figure shows an open document in the JSON editor.




7.2.2.1 About Database Differences in JSON Documents

You can find differences in a JSON document that has been edited from two different connections.

At the time of saving an edited document, you are notified that there are differences between your edited document and the instance of the document in the database. The differences are also highlighted enabling a quick review.

Click **Overwrite** if you prefer to overwrite the changes to the document instance in the database.

Figure 7-1 Differences in the Document Displayed



7.3 About Querying Documents in a Collection

You can search for one or more documents in a collection by using a filter specification or Query-by-Example (QBE).

A QBE is a pattern expressed in JSON. You use it to select, from a collection, the JSON documents whose content matches it, meaning that the condition expressed by the pattern evaluates to true for the content of only those documents. For more information about QBEs and how to use them, see *Overview of SODA QBEs in Oracle Database Introduction to Simple Oracle Document Access (SODA)*.

For a specific collection, enter the QBE string in the JSON editor. For example, to select documents where the name is Mary, enter `{"name":"Mary"}` and then click **Run Query**, as shown in [Figure 7-2](#). The results of the query are displayed in the lower right pane.

The editor offers a comprehensive list of commands available through the Command Palette. To open the Command Palette, press **Ctrl+Shift+P**. For a list of keyboard shortcut keys, see [Keyboard Shortcuts](#).


An error in the query is signified by a red dot in the left gutter and a squiggle line beneath the specific text. When you hover over it, you see a pop-up displaying possible fixes for resolving the error.

You can set editor preferences using the **Preferences** option in the top-right user drop-down list on the header. Some of the available options are Theme (Light, Dark and High contrast dark), Font size and family, Tab size, Word wrap, Ruler, Line numbers and so on.

The icons above the editor are:

- **Format JSON data** : Enables indentation and line feeds for the QBE string.
- **Add Clause**: Adds a formatted template of the \$orderby or \$patch clause to the QBE string.
- **Clear** : Clears the current QBE string.
- **History** : Retrieves previous QBE search strings.

Figure 7-2 QBE String in JSON Editor



The QBE expression must be a valid JSON object and can contain \$query and \$orderby or \$patch clauses. The QBE expression is treated as a \$query clause if there are no clauses defined. Starting with the \$query clause, add the \$orderby or \$patch clause later using the **Add Clause** list. The content is transformed and a template for the \$orderby and \$patch clause is provided. You need to set correct values in the templates. See [Using the In-Context Autocomplete Feature in the JSON Editor](#)

For example, this is a simple filtering query:

```
1 {  
2   "site_admin": true  
3 }
```

After the \$orderby clause is added using Add Clause, you see the following entry:

```
1 {  
2   "$query": {  
3     "site_admin": true  
4   },  
5   "$orderby": {  
6     "$fields": [  
7       {  
8         "path": "",  
9         "datatype": "varchar2",  
10        "order": "asc"  
11      }  
12    ],  
13    "$scalarRequired": false,  
14    "$lax": false  
15  }  
16 }
```

Add Clause ▾
\$orderby
\$patch


After the \$patch clause is added, the query changes to:

```
1 {  
2   "$query": {  
3     "site_admin": true  
4   },  
5   "$patch": [  
6     {  
7       "op": "test"  
8     },  
9     {  
10      "path": "/"  
11    },  
12    {  
13      "value": ""  
14    }  
15  ],  
16   "$orderby": {  
17     "$fields": [  
18       {  
19         "path": "",  
20         "datatype": "varchar2",  
21         "order": "asc"  
22       }  
23     ],  
24     "$scalarRequired": false,  
25     "$lax": false  
26   }  
27 }
```

Add Clause ▾
\$orderby
\$patch

7.3.1 Using the In-Context Autocomplete Feature in the JSON Editor

If you press **Ctrl+Space**, the editor provides you with a context-aware list of options from which you can select and autocomplete at the insertion point.



The following types of information appear in the list:

- **Filter comparison clauses**

A template is available for each clause and it is inserted at the cursor position. For example, when you select the `$between` clause from the autocomplete list, you see the following entry in the editor:



where `age` is the property name and `49` and `70` are values. These are parameters that you can edit and they appear highlighted. You can type or use the autocomplete help to edit the property name. Press the **Tab** key to move to the next parameter.

- **Property names from JSON documents in the collection**

There are two sources for property names from JSON documents:

- If a search index is created with support for data guide, then the JSON data guide from the database dictionary is scanned for property names.
- Property names that are collected from viewed or edited documents are presented. It is possible that the property names are a subset from the whole namespace.

Based on the type of property, the related template is inserted and the cursor is positioned at the expected place for insertion. The following is a list of property types, their templates and the corresponding cursor position for each:

- Object

```
"ShippingInstructions.Address": {
  <cursor_here>
}
```

- Array

"LineItems": [<cursor_here>]

- String

"LineItems[*].Part.Description": "<cursor_here>"

- Number

"PONumber": 0

- Boolean

"site_admin": true

When the cursor is between double quotes (""") and autocompletion is activated, then only the property name is inserted without templates or additional double quotes.

7.4 Creating Indexes for JSON Collections

You can create indexes for JSON collections in the JSON page.

See Also

Indexes for JSON Data in *Oracle Database JSON Developer's Guide*

Open the Indexes Pane

In the JSON left pane, right-click the collection, and select **Indexes**.

The Indexes pane lists the existing indexes for the collection. Select the index row to display more information.

Employees indexes

TABLE ☐ JSON ☐

Name	Type	Schema
emp_name	functional	HR

```

{
  "name": "emp_name",
  "schema": "HR",
  "tableName": "Employees",
  "tableSchemaName": "HR",
  "indexNulls": true,
  "unique": false,
  "lax": false,
  "scalarRequired": false,
  "fields": [
    {
      "path": "name",
      "dataType": "VARCHAR2",
      "maxLength": 2000,
      "order": "ASC"
    }
  ],
  "type": "functional"
}

```

The icons at the top are **Add JSON Index**, **Edit JSON Index**, and **Delete JSON Index**.

The properties of the selected index appear in JSON format below the listed indexes. Select **JSON** from the TABLE - JSON option to view all indexes in JSON presentation.

Create an Index

1. Click the + **New JSON Index** icon. The New Index pane appears.
2. Enter the following fields to create an index:

Name: Type:

☐ Unique ☒ Index Nulls ☐ Path Required

Properties

Enter * to display known properties

☐ Composite index ☐ Advanced

Path	
\$.name	☒
\$.age	☐

- **Name:** Enter a name for the index.
- **Type:** Select the index type from the drop-down list. The different options are Functional, Spatial and Search. Based on the index type selected, the corresponding options appear.

- For a functional type index, the fields to enter are:
 - **Unique:** Select this option to make all indexed values unique.
 - **Index Nulls:** Select this option to use the index in Order By queries.
 - **Path Required:** Select this option if the path must select a scalar value, even a JSON null value.
 - **Properties:** Type the property that you want to index on, or Type * to display all available document properties in the collection. To select a property, select the checkbox in the respective row.

Note

You cannot index properties in arrays.

- **Composite Index:** Select this option if you want to use more than one property.
- **Advanced:** Select this option to change the storage properties of the indexed property. For each property, you can change the type (varchar2, number, date or timestamp), maximum length for indexing (for character properties), and sort order.
- For search index, the options are:
 - **Dataguide off-on:** Select **on** to create JSON data guide for collection.
 - **Text Search off-on:** Select **on** to index all properties in documents to support full-text search based on string equality (every property is treated as string)
 - **Range Search off-on:** Select **on** to support range search when string-range search or temporal search (equality or range) is required.
- Spatial index is used to index GeoJSON geographic data. The selected property should be of GeoJSON type. See [Using GeoJSON Geographic Data](#)

For spatial index, the options are:

- **Path Required:** Select this option if the path must select a value, even if it is a JSON null value.
- **Lax:** Select this option if the targeted field does not need to be present or does not have a GeoJSON geometry object as its value.

Note

You cannot enable **Path Required** and **Lax** at the same time.

3. Click **Create**. A notification is displayed indicating that the index is created and the Indexes pane is populated.

7.5 Creating Relational Views of JSON Documents

You can create relational views of JSON documents.

1. In the JSON page, click  **New Collection View** in the right pane.
2. Enter the following fields:
 - **View Name:** Enter a name for the view.

- **Select columns to add:** Click in this field to view the list of available properties and to select the required properties, adding them one at a time. Alternatively, select **Add All** to add all objects and arrays with all their properties. Select one or several properties and move them up or down to position them in the resulting view. Objects and arrays are moved or deleted together with all their properties.

You can edit the Column Name attribute for each column, and for string types, you can also edit the Length attribute. Precision and scale can be set for numeric properties, for example, (10,2).

Path	Type	Length
name	string	5
address.street	string	13
address.city	string	10
address.zip	number	
address.state	string	2

3. Click **Test Query** to review the DDL statements before creating the view.

Click **Create** to create the view, or click **Definition** to return to the previous screen and make changes to the corresponding attributes.

The created view is now available in the SQL page.


7.6 Viewing the JSON Data Guide Diagram for a Collection

The diagram view displays the JSON data guide for a collection as a hierarchical structure using a format similar to entity-relationship diagrams.

The JSON data guide represents the JSON schema for documents that have a column with JSON content.

In the diagram, arrays are presented as one-to-many relationships, contained objects as one-to-one relationships, and "oneOf" constructs as a box that surrounds possible choices.

Click **Diagram**  in the right pane toolbar to display the JSON data guide diagram for a specific collection.



The icons in the toolbar are Print Diagram, Save to SVG format. Zoom In, Zoom Out, Fit Screen and Actual Size.

See Also

JSON Data Guide in *Oracle Database JSON Developer's Guide*

8

The JavaScript Page

Note

A database user requires the following privileges to use the JavaScript feature:

- CREATE MLE
- EXECUTE DYNAMIC MLE
- EXECUTE ON JAVASCRIPT

With the introduction of Oracle Database Multilingual Engine (MLE), it is possible to run JavaScript directly in Oracle Database. For more information about MLE, see [Introduction to Oracle Database Multilingual Engine](#) in the *Oracle Database JavaScript Developer's Guide*.

The JavaScript feature in Database Actions supports JavaScript code development in Oracle Database 23c and later versions. To use this feature in Database Actions, you must first understand the concepts about MLE, which is explained in detail in the *Oracle Database JavaScript Developer's Guide*.

In the JavaScript interface, you can:

- [Managing MLE Modules](#)
- [Managing MLE Environments](#)
- [Managing MLE Call Specifications](#)
- [Executing JavaScript Code Snippets and Debugging MLE Modules](#)

To navigate to the JavaScript page:


- In the Launchpad page, select the **Development** tab and click **JavaScript**.
- Click **Selector**  to display the navigation menu. Under **Development**, select **JavaScript**.

See Also

[MLE LiveLab \(Includes using Database Actions to create an MLE module, environment and call specification, and debug an MLE module\)](#)

8.1 About the JavaScript Page

The JavaScript page is divided into two parts. The left pane consists of the Navigator and Snippet tabs. The right pane includes the JavaScript code editor at the top and results of the code output at the bottom (this applies only for the Snippet tab in the right pane).



Left Pane

The left pane has the following two tabs:

- **Navigator** to access saved MLE modules, environments and call specifications.
- **Snippets** to access JavaScript code snippets saved in the browser or file system. See [Executing JavaScript Code Snippets and Debugging MLE Modules](#)

In the Navigator tab, you have:

- **Schema and Object type selectors:** Use the drop-down lists to select the schema and filter the results by object type (whether modules, environments or call specifications).
- **Search:** Search for objects in the Navigator tab.
- **Object submenu - Create Object:** Create a new object for the object type.
- **Object submenu - Compile All** (for Environments and Call Specifications): Compile all invalid environment objects or call specification objects.

Code Editor

The JavaScript code editor in the right pane supports autocompletion of text preceding the cursor when using the **Ctrl+Space** key. The autocomplete functionality is triggered for:

- built-in JavaScript modules such as `mle-js-oracledb`.
- export statements of modules in the environment object.
- variable declaration statements exported by the module.

The code editor has two tabs: Editor and Snippet.

Editor Tab

The icons in the Editor tab are:

- **New, Save As, Code Dependencies Diagram:** Create a module, save a module, and view the code dependencies diagram.
- **Open File:** Open a JavaScript file from your file system.

Note

The Open File icon is available only when using a Chromium-based browser in a secure context (HTTPS).

- **Save Module:** Save the module in the database, file system, or in both locations.
- **Import JS Files:** Load multiple JavaScript files from your file system to create them as MLE modules. See [Loading Multiple JavaScript Files to Create MLE Modules](#)

Note

The Import JS Files icon is available only when using a Chromium-based browser in a secure context (HTTPS).

- **Clear:** Remove the contents of the editor.
- **Name:** Enter a name for the MLE module.
- **Environment:** Select the environment for the module.
- **Environment Dependencies Diagram:** Show imported modules and exported functions for each module in the environment in a diagram format.

Snippet Tab

Additionally, the Snippet tab has the following icons:

- **Save:** Save the code snippet that you entered in the editor to the browser or device.
- **Debug Snippet:** Run the debug operation to collect debug information for the selected code snippet.
- **Debug Specification:** Specify the debug specification to use to debug the MLE module.
- **Debug Specification Menu (New, Edit, Delete, Import, Export):** Use the options in this menu to create, edit, or delete a debug specification, and also import or export a debug specification to the file system.

Note

The Import and Export Debug Specification options are available only when using a Chromium-based browser in a secure context (HTTPS).

8.2 Managing MLE Modules

A JavaScript module is a unit of MLE's language code stored in the database as a schema object. You can create, edit and delete JavaScript modules.

Topics:

- [Creating an MLE Module](#)
- [About Context Menu Options for MLE Modules](#)
- [Loading Multiple JavaScript Files to Create MLE Modules](#)
- [Example: Creating an MLE Module](#)

8.2.1 Creating an MLE Module

A JavaScript module can be created in two ways:

- By entering Javascript code into the editor
- By opening the file already available on your local computer

Creating a JavaScript Module by Entering Code in the Editor

To create an MLE module:

1. In the right editor pane, select the **Editor** tab (if not already selected) and enter the JavaScript code.
2. In the **Name** field, enter a name for the module.
3. Select the appropriate environment from the drop-down list.
This is needed when importing or exporting modules already saved to the database and imported in the selected environment.
4. Click **Save** to save in database, file system or both database and file system.

Creating a JavaScript Module by Opening the File

To create an MLE module:

1. In the right pane, select the **Editor** tab, click **Open File** and browse for the file you want. The file contents are displayed in the editor.
2. In the **Name** field, enter a name for the module.
3. Click **Save**.

8.2.2 About Context Menu Options for MLE Modules

In the left Navigator pane, right-click the MLE module name to access the following options:

- **Edit**: Select to edit the attributes of the MLE module.
- **Drop**: Select to remove the MLE module from the database.
- **Dependencies Diagram**: Shows module dependencies in the form of a diagram. The diagram also includes the interlinking of all modules dependent on the specified module. These dependencies are based on the imports and exports in each used module. The details of the exported and imported functions can be seen on the diagram.
- **Create Call Specification**: Opens the Create MLE Call Spec panel. Enables you to create a call specification for the MLE module.
- **Compare (File to Module, Module to File)**: Compares the module version saved in the database (or file system) to the version saved in the file system (or database). When the module is saved in both database and file system (indicated by a round blue overlay icon), use this option to compare the two versions. The differences are highlighted enabling you to quickly identify the changes. Click **Overwrite** to implement the changes needed to synchronize both the database and file system versions of the module.
- **Show Code**: Displays the JavaScript code for the module in a separate panel.
- **Show Metadata**
 - **User Data**: Creates and edits user metadata that can be stored along with the module. The metadata must be in valid JSON syntax.
 - **Code Metadata**: When a module is saved, the metadata related to imports and exports is stored in the database. This metadata is used:
 - * in tracking dependencies on dependency diagrams.
 - * for code completion functionality such as code peek.
 - * for creating an MLE environment.
 - * for creating a call specification based on the module.

Only ES6 modules are supported. There is no support for CommonJS, AMD and UMD modules.

- **Code:** Shows the module code.
- **Generate JSDoc:** Generates the Markdown file based on exports from the module. Data types for function parameters are included if the following exist:
 - JSDoc about function parameters in the code.
 - Call specification defined for the exported function.


8.2.3 Loading Multiple JavaScript Files to Create MLE Modules

You can load multiple JavaScript files to the database as MLE modules. This feature is supported only in Chrome and Chromium-based browsers. A maximum of 150 files can be uploaded in one session.

1. In the JavaScript page, in the right pane, select **Editor** and click **Import JS Files** from the toolbar.
2. In the Import JS Files panel:
 - a. Select the target schema to create the modules.
 - b. Select the environment from the target schema. You can enter a new name to create a new environment.
 - c. Click **Add JS File** and select the files to add.

The files appear in a tabular format in the panel. You can view the module code by clicking the file name. The module name and import name (for the environment) are generated based on the file name but they can be edited.

3. After all the changes are done, click **Create Modules**.




8.2.4 Example: Creating an MLE Module

This example demonstrates the creation of an MLE module.

1. In the JavaScript page, in the left pane, select **Modules** from the object type selector and then click **Object submenu - Create Object**.
2. In the right pane, select the **Editor** tab, and enter the following JavaScript function:

```
export function factorial(num) {  
    if(num < 0) {  
        return -1;  
    } else if (num == 0) {  
        return 1;  
    } else {  
        return(num * factorial(num-1));  
    }  
}
```

3. In the **Name** field, enter **factorial_mod**.



4. Click **Save Module**.

A notification that the module is saved in the database is displayed.

8.3 Managing MLE Environments

MLE environments are schema objects that can be managed in the database. You can create, edit and delete MLE environments using Database Actions.

Creating an MLE Environment

To create an MLE environment:

1. In the left pane, select the **Navigator** tab, select **Object submenu** and click **Create Object**.
The Create MLE Environment panel appears.
2. Select the schema from the list. By default, the current schema is used.
3. In the **Name** field, enter a name for the environment.

4. Select the appropriate **Language Options**. For more information, see [JavaScript Language Options](#) in *JavaScript Developer's Guide*.
5. The Available Modules field automatically populates the modules available for the selected schema. Select one or more modules. Click the **Add Selected Modules** icon or the **Add All Modules** icon to move the modules to the Imported Modules field. You can edit the Import Name field for the module if needed.

When a single module is selected, the **Add module with dependencies** icon becomes available. Click this icon to find environments where the module is used and also to add other modules that the selected module depends on.

To view the code for a module, in the Available Modules field, right-click the module and select **Details**.

6. The DDL pane displays the code for the new MLE environment. Click **Create**.
The MLE environment is saved and appears in the left pane under Environments.

Context Menu Options

In the left pane, right-click the MLE environment to view the following options:

- **Edit**: Select to edit the MLE environment. You can modify the language options, add more modules to import or change the import names for the modules.
- **Drop**: Select to drop the environment from the database.
- **Compile**: Select to compile or recompile the environment.
- **Dependencies Diagram**: Shows the imported modules and exported functions for each module in the environment using a diagram.

See [Example: Creating an MLE Environment](#) for an example of how to create an MLE environment.

8.3.1 Example: Creating an MLE Environment


This example demonstrates the creation of an MLE environment and demonstrates how a mapping is created between the identifier FACTORIAL_MOD and the JavaScript module FACTORIAL_MOD created in [Example: Creating an MLE Module](#).

Create the MLE Environment

1. In the JavaScript page, in the left pane, select **Environments** from the object type selector and then click **Object submenu - Create Object**.
The Create MLE Environment panel appears.
2. In the **Name** field, enter **myfactorialenv**.
3. In the Available Modules field, select **FACTORIAL_MOD**. Click **Add Selected Module** to move the module to the Imported Modules field.

Note

By default, the **Import Name** matches the Module name. You can edit the import name by clicking in the associated column cell and typing in the new name.



4. Click **Create**.

The Output pane displays the underlying source code along with a CREATED notification.

5. Click **Close**.

The new environment is now available in the left Navigator pane. (If you do not see it, click **Refresh** in the left pane.)


Create the MLE module that Imports FACTORIAL_MOD

6. In the JavaScript page, in the left pane, select **Modules** from the object type selector and then click **Object submenu - Create Object**.
7. In the right pane, select the **Editor** tab, and enter the following JavaScript function:

```
import * as h from "FACTORIAL_MOD";

export function factorial_3(){
    console.log('The factorial of 3 is:' + h.factorial(3));
}
```

8. In the **Name** field, enter **mod_imp_fact**.
9. In the **Environment** field, select **MYFACTORIALENV**.




10. Click **Save Module**.

A notification that the module is saved in the database is displayed.

8.4 Managing MLE Call Specifications

Functions exported by an MLE JavaScript module can be published by creating a call specification.

There are two types of call specifications: Module and Inlined (that is JavaScript code embedded in the DDL). Also, call specifications can be created as schema-level functions and procedures or they can be part of package declarations. Accordingly, the Navigator tab on the JavaScript page displays different icons to indicate a call specification in a package, module-based call specification and inlined call specification (with the blue overlay icon).



To create an inlined call specification, in the the **Navigator** tab, select **Object submenu** and click **Inline Call Spec**.

Creating a Call Specification for an MLE Module

To create a call specification for an MLE module:

1. In the left pane, in the **Navigator** tab, select **Object submenu** and click **Create Object**.
The Create MLE Call Spec panel appears.
2. In the **Name** field, enter a name for the call specification.
3. Select the module and environment in the **Module Name** and **Environment Name** fields.
When the module is selected, the functions are automatically populated in the Available Functions box.
4. Select **Procedure** or **Function**. If Function, select the return data type from the drop-down list.
5. Select the function from the list of available functions. When you select the function, the parameters are automatically prefilled in the **Parameters** field. If needed, you can edit the Direction and Type fields. When you click the parameter row, the **Parameter declaration** and **Signature** fields are automatically prefilled.
6. Click **Create**.

See *Creating a Call Specification* in the *Oracle Database JavaScript Developer's Guide* to know more about creating a call specification.

Context Menu Options

In the left pane, right-click the call specification object to view the following options:

- **Edit:** Select to edit the call specification. You can only edit the parameter-related fields.
- **Drop:** Select to remove the call specification from the database.
- **Compile:** Select to compile or recompile the call specification.

See [Example: Creating an MLE Call Specification](#) for an example of how to create a call specification.

8.4.1 Example: Creating an MLE Call Specification

This example demonstrates how to create a call specification for the `mod_imp_fact` module created in [Example: Creating an MLE Environment](#).

1. In the JavaScript page, in the left pane, select **Modules** from the object type selector. From the list of modules displayed, select `mod_imp_fact`.
2. Right-click and select **Create**, and then select **Call Specification**.
The Create MLE Call Spec panel appears.
3. In the **Name** field, enter `factorialspec`.
4. In the **Environment Schema** and **Environment Name** fields, enter your current schema name and select `myfactorialenv`.
5. Select **Function**, and then select **VARCHAR2** from the dropdown list.
6. In the **Available Functions** field, select the `factorial_3` function and the **Signature** field is automatically prefilled.

Create MLE Call Spec

Schema: PDBDBA Name: factorialspec

Properties

Module schema: PDBDBA Module name: MOD_IMP_FACT

Env. schema: PDBDBA Env. name: MYFACTORIALENV

DDL

Procedure: ☒ Function VARCHAR2

Available functions

Name	NAME	DIRECTION	TYPE
factorial_3	No items to display.		

Parameter(s) declaration

Signature *

factorial_3()

Invoker rights clause ☒ DEFINER ☐ CURRENT_USER

Other clauses ☐ PARALLEL ☐ RESULT_CACHE ☐ DETERMINISTIC

Create Close

7. Click **Create**.

The Output pane displays the code for the call specification along with a function compiled notification.

8. Run the function from the SQL Worksheet page and you see the following output:

The factorial of 3 is: 6

8.5 Executing JavaScript Code Snippets and Debugging MLE Modules

You can add, edit, save and execute JavaScript code snippets using the Snippets tab in the right pane of the JavaScript page. Code snippets saved in the browser and file system are listed under the Snippet tab in the left pane.

JavaScript snippets can also be used to test and debug MLE modules using a debug specification. For more information about debug specifications, see [Post-Execution Debugging](#) in the *JavaScript Developer's Guide*.

Note

SQL worksheet supports the execution of JavaScript code in Oracle Database release 21c. These saved JavaScript worksheets will also appear under the Snippets tab in the left pane.

Executing a JavaScript Code Snippet

To execute a code snippet:

1. In the right pane, select the **Snippet** tab and enter the JavaScript snippet. You can do this in one of two ways:
 - Type and enter the code, or copy and paste the code.
 - Click **Open File** in the toolbar to browse and open the file that contains the code.
2. Click the down arrow next to **Snippet** in the toolbar and select **Save As**. Enter a name for the code snippet and save.
3. If required (such as when testing MLE modules), select the preferred environment.
4. Click **Run Snippet**. The Output tab in the lower pane displays the results of the code execution.

Debugging a JavaScript MLE Module

To debug an MLE module:

1. Enter the code snippet, save it and select the environment. See steps 1-3 in [Executing a JavaScript Code Snippet](#).
2. In the Snippet toolbar, from the Debug Specification menu, click **New Debug Specification**.

The Create Debug Specification panel appears.
3. In the Create Debug Specification panel, do the following:
 - a. In the right side of the panel, select the module to debug.

- b. In the left side of the panel, select the JSON template you want: **Snapshot** or **Watch**. For more information about these templates, see [Debugpoint Actions](#) in the *JavaScript Developer's Guide*.
 - c. In the JSON template, the name of the MLE module to debug is specified using the `name` field and the location within the module where the debug information is to be collected is specified using the `line` field. To specify the condition that will trigger the debug action, use the `condition` field. See [Example: Debugging a JavaScript Module](#).
 - d. Click **Create**,
4. In the MLS JS page, select the **Snippet** tab and in the **Debug Specification** field, select the created debug specification.
5. From the Snippet toolbar, click **Debug Snippet**.

The debug information appears under the **Debug Console** tab in the lower pane.

See [Example: Debugging a JavaScript Module](#) for an example of how to debug a JavaScript MLE module.

8.5.1 Example: Debugging a JavaScript Module

This example demonstrates how to debug a JavaScript module.

Create the JavaScript Module

1. In the JavaScript page, in the right pane, select the **Editor** tab and enter the following JavaScript function:

```
export function fib( n ) {  
  if ( n < 2 ) return 1;  
  let n1 = { value : fib( n - 1 ), id : 'n1' };  
  let n2 = { value : fib( n - 2 ), id : 'n2' };  
  return n1.value + n2.value;  
}
```

2. In the **Name** field, enter **fib_mod** and click **Save**.

The created module appears in the left Navigator pane under Modules.

Create the MLE Environment

3. In the Navigator tab, from the object type selector, select **Environments**. Click **Object submenu**.
 4. In the Create MLE Environment panel, do the following:
 - a. In the **Name** field, enter **fib_env**.
 - b. From the **Available Modules** field, select the module **FIB_MOD** and click **Add Selected Modules** to move the module to the **Imported Modules** field.
 - c. Click **Create**.

The new environment appears in the left pane under Environments.

Enter the Code Snippet

5. In the right pane, select the **Snippet** tab and enter the code snippet.

Select the environment **fib_mod** to set the execution context. If the environment does not display, click **Reload Environments** and then expand the dropdown again.

Enter the following code:


```
(async ()=>{  
const aa = await import('FIB_MOD') ;  
console.log(aa.fib(6));  
})();
```

Click the dropdown in the toolbar, select **Save As** and enter a name for the snippet.

6. To execute the snippet, click **Run Snippet** and the output is displayed in the lower right pane under the Output tab.

Create the Debug Specification


7. To create a new debug specification, in the toolbar, click the **Debug Specification Menu** and select **New Debug Specification**.
8. In the Create Debug Specification panel, do the following:
 - a. Select **FIB_MOD** from the **Module** drop-down list.
 - b. In the left side of the panel, select **Snapshot Template**.
 - c. In the JSON document, edit the following:
 - **name**: Enter the module name.
 - **line**: Enter **5** to identify the position in the module code.
 - **condition**: Enter **n > 4**.
 - d. Enter the name **fib_debug** for the debug specification.
 - e. Click **Create**.



9. From the Snippet toolbar, click **Debug Snippet**.

The output appears under the **Debug Console** tab in the lower pane.

If you position the cursor over the debug points, the corresponding position in the module code is displayed.



9

The Scheduling Pages

Note

Available for Oracle Database 12c release 2 and later releases.

Database Actions provides a graphical interface for using the [DBMS_SCHEDULER](#) PL/SQL package to work with Oracle Scheduler objects such as Jobs, Chains, Programs and Schedules. To use the scheduling features, you must first understand the concepts and essential tasks for job scheduling, which are explained in the chapters about Oracle Scheduler concepts and scheduling jobs in the *Oracle Database Administrator's Guide*.

The Scheduling functionality is focused on exploring Oracle Database Scheduler artifacts and provides details for jobs, chains, chain steps and rules, job execution history, programs, schedules, job classes, file watchers, windows and window groups. Depending on the access rights of the logged-in user, information is taken from ALL_* or DBA_* views.

Note

SQL and PL/SQL editors are used in the user interfaces for Jobs and Programs. Some shortcut keys used in these editors are:

- **Alt+F1**: Invokes accessibility options.
- **Ctrl+M**: Alters the Tab key behavior, from inserting tab to going to the next control on the form.

For more keyboard shortcuts, see [Keyboard Shortcuts](#).

To navigate to the Scheduling pages, do either of the following:

- In the Launchpad page, select the **Development** tab and click **Scheduling**.
- Click **Selector**  to display the navigation menu. Under Development, select **Scheduling**.

Topics

- [Overview](#)
- [Jobs](#)
- [Chains](#)
- [Programs](#)
- [Schedules](#)
- [Objects](#)

9.1 Overview

The Overview page provides an overview of the activity in the Scheduling pages. It consists of two parts: Objects and Problem Jobs


The Objects section provides an overview of the total number of running jobs, running chains, programs and schedules. Click an object tile to go to the corresponding page.

Click  **Scheduler Summary** at the top right to view details such as Running Jobs, Current open window, Log history retain, Default timezone, Event expiry time, File watcher count, repeat interval, and mail server settings.

This Problem Jobs section provides information about jobs with problems in their execution. Listed in tabular format are jobs having one of the following statuses: Failed, Broken, Chain_Stalled, Blocked, Retry Scheduled, Resource_Unavailable. The  **Actions** context menu provides links to the history and job details for each problem job. Other actions that you can perform are run, edit, drop, enable and disable job.

Note

If there are no jobs with problems, the Jobs Summary page appears as the initial page for the Scheduling functionality.



	JOB NAME	OWNER	STATE	JOB TYPE	FAILURE COUNT	RUN COUNT	LAST START DATE	NEXT RUN DATE	COMMENTS	
1	EVENT_JOB_SCHED_3	SCHED_DE...	BROKEN	PLSQL_BLOCK	1	5	3/7/2022, 3:01:00 PM	3/7/2022, 3:01:00 PM	test for changes	⋮
2	JOB11	SCHED_DE...	FAILED	PLSQL_BLOCK	2	2	11/17/2021, 12:37:06 PM			⋮
3	JOB_1_PARAM	SCHED_DE...	FAILED	(null)	5	8	2/24/2022, 9:09:16 AM			⋮
4	JOB_ON_PROC_4PAR	SYSTEM	FAILED	STORED_PROCEDURE	1	1	3/11/2022, 10:50:55 AM			⋮
5	JOB_ON_PROC_REM_PROC	SCHED_DE...	FAILED	(null)	1	1	1/24/2022, 11:10:13 AM			⋮
6	JOB_PROG_4_ARGS	SCHED_DE...	FAILED	(null)	8	11	2/23/2022, 12:59:06 PM		job on program with 4 arguments	⋮
7	J_REM_PROC	SCHED_DE...	FAILED	(null)	1	1	3/22/2022, 3:19:40 PM			⋮
8	RUN_MONDAYS	SCHED_DE...	BROKEN	(null)	12	25	3/7/2022, 3:01:00 PM	3/7/2022, 3:01:00 PM		⋮
9										
10										

9.2 Jobs

There are five options available in the Jobs menu: Summary, Running, Forecast, History, Notifications. Each option is described below.


9.2.1 Summary

At the top, the Summary page provides an overview of the total number of Failed, Blocked, Chain stalled, Resource unavailable, Broken, and Retry Scheduled jobs. Click a job tile and the corresponding list of jobs are displayed in tabular format below. Remove the filters to display all scheduled jobs.

To create a job, see [Create or Edit Job](#).

The Actions icon  is available at the end of each job row. Click **Actions** to view the following list of options:

- **Job**
 - **Run:** Runs the specific job.
 - **Edit:** See [Create or Edit Job](#).
 - **Drop:** Drops the specific job.
 - **Job-Enable/Disable:** If enabled is selected, the job is picked up by the Scheduler for processing. The status of the job (enabled or not) is seen in the Job Details page, where the property “enabled” displays TRUE or FALSE.
- **History**
 - **Report:** Provides a history of the job runs in a report format, including log details, status of the run, run duration, errors if any, and so on.
 - **Chart:** Provides a history of the job runs in a visual bar chart representation. Run duration and used CPU time are presented for each execution. Jobs in chains are presented with aggregated information for the whole job. There is an overview scrollbar that enables zoom and scroll by changing the time frame. You can display all details from logs for that time frame for a selected job or for all jobs. Also, a Gantt chart can be shown for all jobs for a selected time frame.



- **Job Forecast:** You can execute the job forecast functionality for a single job (available in the Action menu) or for a set of jobs (available on the toolbar for the table). For a set of jobs, filtering can be used to narrow the set of jobs. Use **Ctrl key+click** to select some jobs and then Job Forecast will run for the selected jobs only. If there are no selected jobs, then forecast will run for all listed jobs. Not every job included in the set is included in the final forecast. Only jobs with a defined calendar (repeat interval) are included. You can define the calendar inline or using schedule, window or window group. After the forecast is done, you can select different zoom levels and filter the results by schema.
- **Job Details:** Displays job attributes, such as action, job class, type, schedule, and so on. Select **JSON** to view the attributes in JSON format. Depending on the job's definition, details about used objects are provided such as used program, PL/SQL code, procedure, procedure dependencies, program arguments, job arguments, job class, schedule,

window, window group together with windows in the group, file watcher, history chart presenting one week of history from last start date.

9.2.2 Running

The Running page displays the list of currently running jobs. The following commands are available: Stop, Edit, History, Job Details.

Scheduling > Running Jobs

Running Jobs

Search by Job Name Filter by

	JOB NAME	OWNER	JOB STYLE	DETACHED	ELAPSED TIME	CPU USED	DESTINATION OWNER	DESTINATION	CREDENTIAL OWNER	CREDENTIAL NAME	LOG ID	
1	A_JOB_ON_CH_N222	SCHED_DEMO_USER	REGULAR	FALSE	P270748M5.5445	null	null	null	null	null	4294967295	
2	CHAIN_JOB_1	SCHED_DEMO_USER	REGULAR	FALSE	P860721H59M32.3295	null	null	null	null	null	4294967296	☐ Stop ☐ Edit ☐ History ☐ Job Details


9.2.3 History

The History page displays log run details for all Scheduler jobs available to the user. You can use the History window and set filters to limit the amount of data. You can filter using delayed jobs by providing a delay interval. Ordering is supported on the grid by clicking the column header.

9.2.4 Forecast


The Forecast page provides the job execution forecast for all available jobs. This operation takes time to finish depending on the number of jobs and forecast interval. Forecast for all available jobs depends on the rights of the connecting user. After execution, you can filter the results based on the schema of the jobs.

The Forecast functionality is also available on the Jobs - Summary page. In that case, it works on the list of available jobs. You can filter the jobs before the forecast functionality is used. If there is a selection of jobs, then the functionality uses only the selected jobs.



9.2.5 History (Gantt Chart)

History represents windows and job execution history in the form of a Gantt chart. The windows activation history is shown on the first row. Jobs are ordered in a descending order based on the maximum used CPU time. Details are shown for each window activation and job execution. The job summary is available when hovering over the label of a job's row.



9.2.6 Notifications

You can create notifications only if an email server is set for the scheduler.

The Notifications page enables you to view, create, edit and delete notifications related to job events. Also, you can see email server details by clicking the **Notifications Email Server** icon at the top right of the page.

For each message, you can specify job, recipient email addresses and sender (or no sender), and you can modify the subject and body of the message and set the filter condition. If multiple recipients and events are provided, then Oracle Scheduler creates a separate notification record for each combination <recipient,event>. You can edit the content of a set of notifications created in this way by using **Edit Aggregated**, which appears in the context menu of each notification. **Edit** edits the content of a single notification.

Notifications can be filtered by job name, job owners, recipients and events. You can select and remove some notifications by clicking the **Remove Notifications** icon. If there are no selected notifications (use **Ctrl+click** to deselect a single notification), then the Remove Notifications dialog appears enabling removal of all notifications for the selected job, or job and recipients, or job and events, or job and events and recipients.

Remove Notifications

Properties

SQL

Schema

SCHED_DEMO_USER

Job Name *

DW_ETL_JOB

Select Recipients

	☐	Recipient
1	☐	joel@abc.com
2	☐	smith@abc.com

Select Events

	☐	Event
1	☐	JOB_DISABLED
2	☐	JOB_FAILED
3	☐	JOB_SCH_LIM_REACHED

?

Remove

Cancel

9.2.7 Create or Edit Job

This section describes how to create a new Oracle Scheduler job or edit an existing job.

To create a job, Database Actions internally uses the `DBMS_SCHEDULER.CREATE_JOB` procedure, which is documented in *Oracle Database PL/SQL Packages and Types Reference*.

1. In the Jobs page, at the top right, click **Create Job**.

2. In Job Properties, enter the following fields:

Details

- **Enabled:** If this option is specified, validity checks are made and the job is created enabled if all the checks are successful. If this option is not specified, the job is not created enabled.
- **Name:** Name of the job.
- **Description:** Optional text string that can be used to describe the job.
- **Type:** Type of object to be executed by the job: PL/SQL Block, Chain, Stored Procedure, Named Program, or Script. Additional controls appear for Chain, Stored Procedure and Named Program enabling you to select related objects.

For Stored Procedure, only procedures with IN parameters are listed, procedures with IN OUT or OUT parameters are not permitted.

Schema level and package level procedures are listed for the selected schema but you can directly type the procedure name (or package_name.procedure_name) in the field.

- **Class:** Name of the job class to which this job belongs.

Execution Mode

- **Mode:** When to execute the job: **Immediate** (immediately on creation, and once only), **Once** (once, at a specified time), **Repeating**, **Queue**, **File Watcher**, **Schedule** (using a named schedule object), **Window** and **Window group**. If you specify anything other than Immediate, you are prompted for additional information.
For Repeating, you can manually define the repeat interval or click the pencil icon to select the date, frequency, weekday, and interval values. Some clauses of Oracle calendar syntax (include, exclude, intersect, periods and by period) are not supported and a warning is displayed when the edit icon is clicked.

Destination

- **Local** (local system), **Remote** (the database destination for a remote database job, or external destination for a remote external job), or **Multiple** (the job runs on all destinations associated with the provided destination group). Depending on what destination you selected for the job, select the local credential, the remote credential and destination, or the destination group.

Properties

- **Auto Drop:** Determines whether the job is to be automatically dropped after it has completed or has been automatically disabled.
- **Restart on failure:** Determines whether the job can be restarted in case of failure.
- **Restart on recovery:** Determines whether to restart the job in case of database failure.
- **Store Output:** If enabled, then for job runs that are logged, all job output and error messages are stored in the *_JOB_RUN_DETAILS views. If disabled, then the output and messages are not stored.
- **Follow Default Time Zone:** Determines whether if the job start date is null, then when the default time zone scheduler attribute is changed, the Scheduler recomputes the next run date and time for this job so that it is in accordance with the new time zone.
- **Allow Runs in Restricted Mode:** If enabled, the job is permitted to run when the database is in restricted mode, provided that the job owner is permitted to log in during this mode.

- **Stop on Window Close:** If the schedule of a job is a window or a window group, enabling this option causes the job to stop once the associated window is closed, and disabling causes the job to continue after the window closes. (Note that if the job is allowed to continue, its resource allocation will probably change because closing a window generally also implies a change in resource plans.)
- **Instance Stickiness:** This attribute should only be used for a database running in an Oracle Real Application Clusters (Oracle RAC) environment. By default, it is enabled. Jobs start running on the instance with the lightest load and the Scheduler thereafter attempts to run on the instance that it last ran on. If that instance is either down or so overloaded that it does not start new jobs for a significant period of time, another instance runs the job. If the interval between runs is large, `instance_stickiness` is ignored and the job is handled as if it were a non-sticky job. If `instance_stickiness` is disabled, each instance of the job runs on the first instance available.
- **Parallel Instances:** For an event-based job, determines what happens if an event is raised and the event-based job that processes that event is already running. If disabled, it causes the new event to be ignored. If enabled, it causes an instance of the job to be started for every instance of the event, and each job instance is a lightweight job so multiple instances of the same event-based job can run in parallel.
- **Job Style:** Style of the job being created: **REGULAR** (regular job) or **LIGHTWEIGHT** (lightweight job). A lightweight must reference a program object. Use lightweight jobs when you have many short-duration jobs that run frequently. Under certain circumstances, using lightweight jobs can deliver a small performance gain.
- **Job Priority:** The priority of this job relative to other jobs in the same class as this job. If multiple jobs within a class are scheduled to be executed at the same time, the job priority determines the order in which jobs from that class are picked up for execution by the job coordinator. It can be a value from 1 through 5, with 1 being the first to be picked up for job execution.
- **Logging Level:** Determines how much information is logged:
`DBMS_SCHEDULER.LOGGING_OFF` (no logging),
`DBMS_SCHEDULER.LOGGING_FAILED_RUNS` (only jobs that failed, with the reason for failure), `DBMS_SCHEDULER.LOGGING_RUNS` (all runs of each job in this class), or `DBMS_SCHEDULER.LOGGING_FULL` (all operations performed on all jobs).
 However, if the job class has a higher (more detailed) logging level than the level specified for the job, the job class logging level is used.
- **Instance ID:** In an Oracle Real Application Clusters environment., the instance ID of the instance that the job must run on.
- **Max Runs:** The maximum number of consecutive scheduled runs of the job.
- **Max Failures:** The number of times a job can fail on consecutive scheduled runs before it is automatically disabled.
- **Raise Events:** Determines at what stages of the job execution to raise events.
- **Max Run Duration:** Maximum amount of time that the job should be allowed to run. Its data type is `INTERVAL DAY TO SECOND`. If this attribute is set to a nonzero and non-null value, and job duration exceeds this value, the Scheduler raises an event of type `JOB_OVER_MAX_DUR`. It is then up to your event handler to decide whether or not to allow the job to continue.
- **Schedule Limit:** Maximum delay time between scheduled and actual job start before a program run is canceled.
- **Reset to Defaults:** Resets all properties to their default values.

NLS

Enables you to set the NLS-related property required for this job. To enter a value, type in to the related field for the Value column.

For a new job, parameters are taken from the database session.

3. In the **DDL** pane, you can review and save the SQL statements that are generated when creating or editing the job.
 - For a new job, click **CREATE** to view the generated DDL statements.
 - When you edit a job, click **UPDATE** to view the generated ALTER statements.

When you are finished, click **Apply**.

The **Output** pane displays the results of the DDL commands. If there are any errors, go to the corresponding pane, fix the errors, and run the commands again. You can save the content to a file.

9.3 Chains


There are four options available in the Chains menu: Summary, Running Chains, Steps, Rules. Each option is described below.

Summary

The Summary page displays all the chains owned by the current user. The details are presented in a tabular format. At the end of each row is the Actions icon . Click **Actions** to view the list of options for the chain. Click a chain in the table and details about steps and rules appear below the Chains table. There is a **Drop Step** or **Drop Rule** option available in the context menu for a step or rule.

The actions available for a chain are:

- **Run:** Runs the specific chain.
- **Add Step:** Creates a new step for the selected chain.
- **Add Rule:** Creates a new rule for the selected chain.
- **Edit:** Edits the chain properties.
- **Drop:** Drops the chain
- **Chain-Enable/Disable:** Enables or disables the chain.
- **Show Diagram:** Displays the job steps of a chain in a visual diagram format.




Double-click a step or rule in the diagram to see more details. You can directly create steps and rules on the diagram using the **Add Step** and **Add Rule** icons.

The following dependencies are shown on each step:


- Event schedule: The name of the schedule.
- Event queue: The name of the queue and condition.
- Program using procedure or PL/SQL block: Tables, views, procedures functions and packages used - two levels of dependencies.

You can also create a rule by connecting two steps:

- Click on a step and the link icon appears.
- Click the link icon and drag to the next step. The Rule dialog appears where you can provide details.



There are validation rules that are applied to chain details and if problems are found, an icon appears on the menu bar that enables you to get details about the rules violated.



The following rules are applied:

- Usage of missing steps in chain rule
 - Steps without start: There is no rule that starts the step.
 - Steps without transition to another step or to the end: The step does not appear in the condition part of any rule.
 - Steps with disabled Program: of type “Warning” – show steps using disabled programs
 - Steps depending on invalid objects: of type “Warning” – show steps that use program dependent on invalid objects.
- **Analyze Chain:** Displays output from the `DBMS_SCHEDULER.ANALYZE_CHAIN` procedure in text format. Lists the type of steps, dependencies and transitions between them based on defined rules.

When a chain row is selected, the lists with steps and rules appear below. The Edit and Drop commands are available in the Actions menu for Steps and Rules.

Running Chains

The Running Chains page displays the execution of the steps of the currently running chains.

The following commands are available for each row in the report:

- For Chain, the commands are Edit, Drop, Enable, Disable, and Evaluate Chain.
- Evaluate Chain** invokes the `DBMS_SCHEDULER.EVALUATE_RUNNING_CHAIN` procedure. It forces the reevaluation of the rules of a running chain to trigger any rules for which the conditions have been satisfied.
- Other chain-related command are Show Diagram, Analyze Chain, Step Details, and Edit Step.

Steps

The Steps page provides details of each step in all chains.

Rules

The Rules page provides details of each rule in all chains.

9.3.1 Create or Edit Chain

This section describes how to create or edit a chain.

To create a chain, Database Actions internally uses the DBMS_SCHEDULER.CREATE_CHAIN procedure.

1. In the Chains page, at the top right, click **Create Chain**.
2. In Properties, enter the following fields:
 - **Name:** Name to assign to the chain.
 - **Evaluation Interval:** **NULL** reevaluates of the rules of a running chain only when the job starts and when a step completes. A non-NULL value causes rule evaluations to also occur periodically at the specified interval.
 - **Enabled:** Enables the chain. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the chain is created.)
3. Click **Create** to create the chain.

Click **Create and Open** to create the chain and open in Diagram mode.

9.3.2 Add or Edit Step

This section describes how to add a step for a chain.

1. In the context menu for a chain, select **Add Step**.
2. In the Create Step slider, enter the following fields:
 - **Step Name:** Name of the step.
 - **Type:** Select **Program**, **Chain**, **Event** or **Queue** to run during the step. Additionally, select the name of the program, chain, event or queue in the corresponding field.
 - **Pause:** Select this option to pause the step.
 - **Skip:** Select this option to skip the step.
 - **Restart on Failure:** Select this option to restart the step after database failure.
 - **Restart on Recovery:** Select this option to restart the step after database recovery.
3. Click **Create**.

The new step is displayed for the chain in the main Chains page.

9.3.3 Add or Edit Rule

This section describes how to add a rule for a chain.

1. In the context menu for a chain, select **Add Rule**.
2. In the Create Rule slider, enter the following fields:
 - **Rule Name:** Name of the rule being created.

- **Condition:** Enter the condition to evaluate.
- **Action:** Action to perform if condition is met.
- **Comments:** Optional comments describing the rule.

See Also

[DEFINE_CHAIN_RULE Procedure](#) in *Oracle Database PL/SQL Packages and Types Reference*

9.4 Programs

The Programs page displays the list of created programs.

The Actions menu for a program consists of the following options: Edit, Drop, Enable or Disable the program and View Program Details, including used procedure and procedure dependencies.

Programs							RANDOM_EXEC_TIME	
Search by Program Name Filter by							Program Properties	
	Detached	Schedule Limit	Priority	Weight	Max Runs	Max Failures	Max Run Duration	Procedure PL/SQL
1	FALSE	(null)	3	1	(null)	(null)	(null)	Procedure Dependencies
2	FALSE	(null)	3	1	(null)	(null)	(null)	
3	FALSE	(null)	3	1	(null)	(null)	(null)	
4	FALSE	(null)	3	1	(null)	(null)	(null)	
5	FALSE	(null)	3	1	(null)	(null)	(null)	
6	FALSE	(null)	3	1	(null)	(null)	(null)	
7	FALSE	(null)	3	1	(null)	(null)	(null)	
8	FALSE	(null)	3	1	(null)	(null)	(null)	
9	FALSE	(null)	3	1	(null)	(null)	(null)	
10	FALSE	(null)	3	1	(null)	(null)	(null)	
11	FALSE	(null)	3	1	(null)	(null)	(null)	
12	FALSE	(null)	3	1	(null)	(null)	(null)	
13	FALSE	(null)	3	1	(null)	(null)	(null)	
14	FALSE	(null)	3	1	(null)	(null)	(null)	
15	FALSE	(null)	3	1	(null)	(null)	(null)	
16	FALSE	(null)	3	1	(null)	(null)	(null)	
17	FALSE	(null)	3	1	(null)	(null)	(null)	
18	FALSE	(null)	3	1	(null)	(null)	(null)	
19	FALSE	(null)	3	1	(null)	(null)	(null)	

```

PROCEDURE RANDOM_EXEC AS
  CURR_TIMES TIMESTAMP WITH TIME ZONE := SYSTIMESTAMP;
  CURR_TIMES NUMBER := DBMS_UTILITY.GET_TIME;
  END_TIME TIMESTAMP;
  GEN NUMBER(2);
BEGIN
  END_TIME := CURR_TIMES + INTERVAL '3' SECOND;
  GEN := DBMS_RANDOM.VALUE(1, 10);
  IF GEN IN (1, 2, 3) THEN
    END_TIME := CURR_TIMES + INTERVAL '15' SECOND;
  ELSEIF GEN IN (4, 5, 6) THEN
    END_TIME := CURR_TIMES + INTERVAL '5' SECOND;
  END IF;
  DBMS_OUTPUT.PUT_LINE(GEN);
  LOOP
    EXIT WHEN CURRENT_TIMESTAMP >= END_TIME;
  END LOOP;
END RANDOM_EXEC;

```

9.4.1 Create or Edit Program

This section describes how to create or edit an Oracle Scheduler program.

To create a program, Database Actions internally uses the `DBMS_SCHEDULER.CREATE_PROGRAM` procedure.

1. In the Programs page, at the top right, click **Create Program**.
2. In Program Properties, enter the following fields:

Details tab

- **Name:** Name of the program. The name has to be unique in the SQL namespace. For example, a program cannot have the same name as a table in a schema.
- **Enabled:** If this option is specified, validity checks will be made and the program will be created enabled if all the checks are successful. If this option is not specified, the program is not created enabled.
- **Description:** Optional text string that can be used to describe the program.

- **Type**
 - **PL/SQL Block:** The program is a PL/SQL block. Job or program arguments are not supported when the job or program type is PLSQL_BLOCK. In this case, the number of arguments must be 0. Enter or paste in the complete PL/SQL code, or edit the existing code.
 - **Stored Procedure:** The program is a PL/SQL or Java stored procedure, or an external C subprogram. Only procedures, not functions with return values, are supported. PL/SQL procedures with IN OUT or OUT arguments are not supported.

Schema: Schema of the stored procedure. If not specified, the schema of the job is assumed.

Procedure: Name of the stored procedure.

Arguments: For each argument, name, data type, default value, and whether it is an input, output, or input/output argument.
 - **Script:** The program is a SQL Script (SQL*Plus statements), Backup Script (RMAN commands), or External Script (operating system commands). Enter or paste the script text in the box.

Properties tab


Enables you to set program properties. For most properties the default is null, but you can check the box to specify a value.

- **Detached:** Enabled if the program is a detached job. Use a detached job to start a script or application that runs in a separate process, independently and asynchronously to the Scheduler. A detached job typically starts another process and then exits. Upon exit (when the job action is completed) a detached job remains in the running state. The running state indicates that the asynchronous process that the job started is still active. When the asynchronous process finishes its work, it must connect to the database and call DBMS_SCHEDULER.END_DETACHED_JOB_RUN, which ends the job.
 - **Max Runs:** Maximum number of runs before the program is marked as completed.
 - **Max Failures:** Maximum number of failures tolerated before the program is marked as broken.
 - **Max Run Duration:** Maximum run duration of the program.
 - **Schedule Limit:** Maximum delay time between scheduled and actual job start before a program run is canceled.
3. In the **DDL** pane, you can review and save the SQL statements that are generated.
- For a new program, click **CREATE** to view the generated DDL statements.
 - When you edit a program, click **UPDATE** to view the generated ALTER statements.
- When you are finished, click **Apply**.
4. The **Output** pane displays the results of the DDL commands. If there are any errors, go to the corresponding pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

9.5 Schedules

The Schedules page displays the list of available schedules.

The Actions menu for a schedule has the following options: Viewing schedule details, Editing or Dropping the schedule.



9.5.1 Create or Edit Schedule

This section describes how to create or edit Oracle Scheduler schedule.

To create a schedule, Database Actions internally uses the DBMS_SCHEDULER.CREATE_SCHEDULE procedure.

1. In the Schedules page, at the top right, click **Create Schedule**.
2. In Program Properties, enter the following fields:

Properties tab

- **Name:** Name of the schedule. The name has to be unique in the SQL namespace. For example, a schedule cannot have the same name as a table in a schema.
 - **Description:** Optional text string that can be used to describe the schedule.
 - **Mode:** Specify when jobs that use this schedule are to run: **Repeating** (specify the repeat interval, start date, and end date), **Queue** (specify the queue name, agent, condition, start date, and end date), or **File Watcher** (specify the file watcher object name, condition, start date, and end date).
3. In the **DDL** pane, you can review and save the SQL statements that are generated.
 - For a new schedule, click **CREATE** to view the generated DDL statements.
 - When you edit a schedule, click **UPDATE** to view the generated ALTER statements.

When you are finished, click **Apply**.

4. The **Output** pane displays the results of the DDL commands. If there are any errors, go to the corresponding pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

9.6 Objects

The Objects menu in Scheduling provides links to the following four pages: Job Classes, File Watchers, Windows and Window Groups

Job Classes

A job class is an Oracle Scheduler object that enables the Scheduler administrator to group jobs for logical purposes, such as to assign the same set of attribute values to member jobs, to set service affinity for member jobs, to set resource allocation for member jobs, or to group jobs for prioritization.

Note

Creating a job class requires the `MANAGE SCHEDULER` system privilege. Access to `SYS.DBA_RSRC_CONSUMER_GROUPS` and `SYS.DBA_SERVICES` dictionary views is required in order UI to provide all details and choices.

The Job Classes page displays the job classes created in a tabular format. You can perform the following actions:

- **Create:** To create a job class, see [Create Job Class](#).
- **Edit:** To edit a job class, select **Edit** from the Actions menu at the end of a row. For a description of the fields, see [Create Job Class](#).
- **Drop:** To drop a job class, select **Drop** from the Actions menu.

File Watchers

A file watcher is an Oracle Scheduler object that defines the location, name, and other properties of a file whose arrival on a system causes the Scheduler to start a job. You create a file watcher and then create any number of event-based jobs or event schedules that reference the file watcher. When the file watcher detects the arrival of the designated file, it raises a file arrival event. The job started by the file arrival event can retrieve the event message to learn about the newly arrived file.

The File Watchers page displays the file watchers created in a tabular format. You can perform the following actions:

- **Create:** To create a file watcher, see [Create File Watcher](#).
- **Edit:** To edit a file watcher, select **Edit** from the Actions menu at the end of a row. For a description of the fields, see [Create File Watcher](#).
- **Drop:** To drop a file watcher, select **Drop** from the Actions menu.
- **Show DDL:** Displays the Data Definition Language statements for the file watcher. To show DDL, select this option from the Actions menu.

Windows

A window is an Oracle Scheduler object that can be used to automatically start jobs or to change resource allocation among jobs during various time periods of the day, week, and so on.

Note

Creating a window object requires the `MANAGE SCHEDULER` system privilege. Access to `SYS.DBA_RSRC_PLANS` dictionary view is required to provide all details and choices.

The Windows page displays windows created in a tabular format. You can perform the following actions:

- **Create:** To create a window, see [Create Window](#).
- **Edit:** To edit a window, select **Edit** from the Actions menu at the end of a row. For a description of the fields, see [Create Window](#).
- **Drop:** To drop a window, select **Drop** from the Actions menu.

- **Show DDL:** Displays the Data Definition Language statements for the window. To show DDL, select this option from the Actions menu.

Window Groups

A window group is an Oracle Scheduler object that is a list of Oracle Scheduler Windows. Scheduler jobs that are scheduled to be run in a window group will be activated in that time span and using that resource plan for all windows in the group

The Window Groups page displays windows created in a tabular format. You can perform the following actions:

- **Create:** To create a window group, see [Create Window Group](#).
- **Edit:** To edit a window group, select **Edit** from the Actions menu at the end of a row. For a description of the fields, see [Create Window Group](#).
- **Drop:** To drop a window group, select **Drop** from the Actions menu.
- **Show DDL:** Displays the Data Definition Language statements for the window group. To show DDL, select this option from the Actions menu.

9.6.1 Create Job Class

To create a job class:

1. In the Job Class page, click **Create Job Class**.
2. Enter the following fields:
 - **Name:** Name of the job class.
 - **Description:** Optional text string that can be used to describe the job class.
 - **Logging Level:** Specifies how much information is written to the job log:
 - **RUNS:** Detailed information for all runs of each job in this class.
 - **FULL:** RUNS plus all operations performed on all jobs in this class.
 - **OFF:** No logging.
 - **Log Retention Period (days):** Number of days that job log entries for jobs in this class are retained. The range of valid values is 0 through 999. If set to 0, no history is kept. If NULL (the default), retention days are set by the log_history Scheduler attribute.
 - **Service Name:** The database service that the jobs in this class will have affinity to. If no service is specified, the job class will belong to the default service, which means it will have no service affinity and any one of the database instances within the cluster might run the job.
 - **Resource Consumer Group:** Resource consumer group this class is associated with. If no resource consumer group is specified, the job class is associated with the default resource consumer group.
3. Click **Create**.

9.6.2 Create File Watcher

To create a file watcher:

1. In the File Watchers page, click **Create File Watcher**.
2. Enter the following fields:

- **Enabled:** Enables the file watcher. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the file watcher is created.)
- **Name:** Name of the file watcher.
- **Description:** Optional descriptive text.
- **Destination:** Name of an external destination. You create an external destination by registering a remote Scheduler agent with the database. The view ALL_SCHEDULER_EXTERNAL_DESTS lists valid external destination names. If this parameter is null, the file watcher is created on the local host.
- **Credential Name:** Name of a valid credential object. The file watcher uses the credential to authenticate itself with the host operating system to access the watched-for file. The file watcher owner must have EXECUTE privileges on the credential.
- **Directory Path:** Directory in which the file is expected to arrive. The single wildcard '?' at the beginning of the path denotes the Oracle home path. For example, '?/rdbms/log' denotes the rdbms/log subdirectory of the Oracle home directory.
- **File Name:** Name of the file to look for. Two wildcards are permitted anywhere in the file name: '?' denotes any single character, and '*' denotes zero or more characters.
- **Min File Size:** Minimum size in bytes that the file must be before the file watcher considers the file found.
- **Steady State Duration:** Minimum time interval that the file must remain unchanged before the file watcher considers the file found. Cannot exceed one hour. If null, an internal value is used. The minimum value is 10 seconds. Oracle recommends similar steady state duration values for all file watchers for efficient file watcher job operation. Also, the repeat interval of the file watcher schedule must be equal or greater than the steady state duration value.

3. Click **Create**.

9.6.3 Create Window

To create a window:

1. In the Windows page, click **Create Window**.
2. Enter the following fields:
 - **Name:** Name of the window.
 - **Description:** Optional descriptive text.
 - **Enabled:** Enables the window. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the window is created.)
 - **Resource Plan:** The resource plan that automatically activates when the window opens. When the window closes, the system switches to the appropriate resource plan, which is usually the plan that was in effect before the window opened, but can also be the plan of a different window. Only one resource plan can be associated with a window. If null, the resource plan in effect when the window opens stays in effect for the duration of the window. If an empty string, the resource manager is disabled for the duration of the window. If the window is open and the resource plan is dropped, then the resource allocation for the duration of the window is not affected.
 - **Duration:** The length of time that the window stays open. Can range from one minute to 99 days.

- **Priority:** Relevant when two windows overlap. Because only one window can be in effect at one time, the window priority determines which window opens. The two possible values for this attribute are HIGH and LOW. A high priority window has precedence over a low priority window, therefore, the low priority window does not open if it overlaps a high priority window.
- **Execution Mode:** Repeating or Schedule:
 - For Repeating, specify the **Repeat Interval** and optionally the **Start Date** and **End Date**.
 - For Schedule, select the name of the **Schedule** to be used.

3. Click **Create**.

9.6.4 Create Window Group

To create a window group:

1. In the Window Groups page, click **Create Window Group**.
2. Enter the following fields:
 - **Name:** Name of the window group.
 - **Description:** Optional comments.
 - **Enabled:** Enables the window group. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the window group is created.)
 - **Available Windows:** Lists all Scheduler windows.
 - **Selected Windows:** List windows to be added to the window group. Use the arrow icons to move selected windows or all windows from one list to the other.
3. Click **Create**.

Part II

Vector Database

This part describes how to use the Vector Database Console to manage models, create vector tables, load data, and perform semantic searches.

- [The Vector Database Page](#)

The Vector Database Page

The Vector Database page provides a centralized way to manage vector-specific operations such as loading embedding models, creating tables with vector data, and performing semantic search in a convenient query playground.

Topics:

- [Install the Python SDK](#)
- [Use the Vector Database Console](#)
- [Manage Models](#)
- [Manage Vector Tables](#)
- [Manage Data](#)
- [Search Data](#)
- [Manage Indexes](#)
- [Manage Vector Database Users](#)
- [Configure an OAuth Client](#)

Install the Python SDK

The Python SDK, `oracle-vecdb`, is distributed using Python's package repository (PyPI) and can be installed using `pip`.

First ensure that Python is installed on your system. The versions supported with `oracle-vecdb` are Python 3.10 and later. You can check your version using the following command:

```
python --version
```

Note

On some platforms, the Python executable may be called `python3` instead of `python`.

Run the following command to install or upgrade oracle-vecdb:

```
python -m pip install oracle-vecdb --upgrade
```

Using `python -m pip` ensures that the package is installed for the same Python interpreter that you are running.

If you are not using a virtual environment and do not have permission to install packages system-wide, you can choose to append the `--user` option to the end of the install command to indicate a user-level installation.

You can verify the installation by importing the `oracle_vecdb` package in Python:

```
python -c "import oracle_vecdb"
```

If no error is raised, the installation was successful.

Create a Client and Perform Simple Queries

Once you have installed the Python SDK, you can follow these example to first create a client and then create a table, upsert vectors, and perform a semantic search with filters. Note that you will need to input your own `rest_url` and `access_token` (or username and password) values to create the client.

1. Configure the client:

```
from oracle_vecdb import OracleVecDB, Configuration

config = Configuration(
    rest_url="https://<host>:<port>/ords/<schema>/_/_db-api/stable/vecdb/",
    # choose one auth method
    access_token="<bearer-token>",
    # or username="<user>", password="<pass>",
)

vecdb = OracleVecDB(config)
```

Note

Proxies and retries can be configured using `Configuration`.

2. Create an integrated embedding table.

Create a table that generates embeddings from text stored in metadata. The configured model must already be available in Oracle AI Database.

```
vecdb.create_vector_table(
    name="demo",
    table_params={"auto_generate_id": True},
    embed_params={
        "model": "all_MiniLM_L12_v2",
        "embed_metadata_jsonpath": "content",
    },
)
```

3. Load model-backed records.

When integrated embedding is configured, provide text in the metadata field selected by `embed_metadata_jsonpath` (in this case, `"content"`). The database generates the vector during upsert.

```
vecdb.upsert_vectors(
    table_name="demo",
    vectors=[
        {
            "metadata": {
                "title": "Comedy movie review",
                "content": "A lighthearted comedy with fast-paced jokes.",
                "genre": "comedy",
            }
        },
        {
            "metadata": {
                "title": "Drama movie review",
                "content": "An emotional family drama with strong performances.",
                "genre": "drama",
            }
        }
    ],
)
```

4. Run a text query with filtering.

A text query uses the table's configured embedding model to generate the query vector.

```
results = vecdb.query(
    table_name="demo",
    query_by={"text": "family drama"},
    filters={"genre": {"$eq": "drama"}},
    top_k=1,
)

for index in range(len(results)):
    item = results[index]
    row = item if isinstance(item, dict) else item.model_dump()
    print(row["id"], row["distance"], row["metadata"])
```

Create a Table Bringing Your Own Vectors

For precomputed embeddings, omit the `embed_params` property from the vector table creation and provide `dense_vector` values in each record.

```
vecdb.create_vector_table(name="demo_byov")
vecdb.upsert_vectors(
    table_name="demo_byov",
    vectors=[
        {"id": "1", "dense_vector": [0.1, 0.1], "metadata": {"genre":
"comedy"}},
        {"id": "2", "dense_vector": [0.2, 0.2], "metadata": {"genre":
"drama"}},
    ],
)
```

```

)

results = vecdb.query(
    table_name="demo_byov",
    query_by={"vector": [0.15, 0.1]},
    filters={"genre": {"$eq": "drama"}},
    top_k=1,
)

for index in range(len(results)):
    item = results[index]
    row = item if isinstance(item, dict) else item.model_dump()
    print(row["metadata"]["genre"])

```

Create Indexes and Perform Tuning

- Delay index creation.

Create the table first and build its index explicitly when the data-loading workflow is complete.

```

vecdb.create_vector_table(
    name="demo_manual",
    index_params={"vector_index_params": {"auto_index": False}},
)

```

```

vecdb.create_index(table_name="demo_manual")

```

- Create an Hierarchical Navigable Small World (HNSW) index.

Use INMEMORY GRAPH as the value for organization for an HNSW-style vector index.

```

vecdb.create_vector_table(
    name="demo_hnsw",
    index_params={
        "vector_index_params": {
            "organization": "INMEMORY GRAPH",
            "advanced_params": {
                "neighbors": 32,
                "efConstruction": 200,
            },
        },
    },
)

```

- Query-time HNSW tuning.

Tune HNSW search with advanced_options without rebuilding the index.

```

results = vecdb.query(
    table_name="demo",
    query_by={"text": "family drama"},
    filters={"genre": {"$eq": "drama"}},
    top_k=1,
    advanced_options={
        "idx_parameters": {
            "efsearch": 64,

```

```
}  
},  
)
```

For more in depth examples that demonstrate how to use Vector Database Console, see the following resources on GitHub:

- [Sample Applications](#)
- [Sample Notebooks](#)

Use the Vector Database Console

The Vector Database Console provides a low-code user interface that can be used to manage resources and perform inference on your Vector Database.

Access the Vector Database Console using the following steps:

- From your Vector Database instance:
 1. Click the **Database actions** dropdown menu, and select **View all database actions**.
 2. Select the **Vector Database** tab, and click **Open**.

Note

You are initially logged in as an Administrator user. To create a Vector database user, see [Create a New Vector Database User](#). After creating the user, log out of **Database Actions**, and then log back in as the newly created vector database user.

- From the Database Actions Launchpad page:
 1. Click **Selector**  to display the navigation menu.
 2. Under **Vector Database**, select **Vector Database**.

The first page displayed by the Vector Database Console is the **Overview** page, which provides links to getting started tasks, including choosing and loading a model, creating a vector table, upserting data, and performing semantic search.

The following pages are accessible from the left hand navigation menu:

- **Overview:** Find links to get started using your Vector Database along with an account of the current number of tables and loaded models.
- **Tables:** View, manage, and create tables in your Vector Database.
- **Query playground:** Perform text and vector searches against your Vector Database. Advanced options allow you to customize queries by specifying Top K results and adding search filters.
- **Models:** View the list of pre-configured models that you can use to vectorize your data. Choose which model(s) you would like to use and load them into your database.
- **Sample apps:** A GitHub repository with sample applications that demonstrate the use and benefits of the Vector Database.

Topics:

- [Manage Models Using the Vector Database Console](#)

- [Manage Vector Tables Using the Vector Database Console](#)
- [Use the Query Playground](#)

Manage Models Using the Vector Database Console

You can use the vector database console to perform management tasks, such as loading and dropping models.

Once you are in the vector database console, select the **Models** tab to view and manage the available pre-configured models.

- [Load a Model Using the Vector Database Console](#)
- [Remove a Model Using the Vector Database Console](#)

Load a Model Using the Vector Database Console

Follow these steps to load a pre-configured model using the vector database console.

1. From the vector database console, select the **Models** tab.
This will take you to a list of available pre-configured models that you can use to vectorize your data.
2. Select the three dots  icon in the tile of your desired model, then click the **Load** option.
You will be prompted to confirm your choice. Once selected, the model will display a status of **Loading** until it becomes **Available**.

Remove a Model Using the Vector Database Console

Follow these steps to remove a model that you have previously loaded using the vector database console.

1. From the vector database console, select the **Models** tab.
2. Click the three dots  in the row of the model you want to remove. Then select **Remove model**.
The **Remove model** option appears for any model that has already been loaded and has a status of **Available**.
3. Confirm that you would like to remove the model by again selecting **Remove model** when prompted.
 - If a table exists that is dependent upon the selected model, an error message will be displayed that informs you of the impacted tables. In this case, either select **Close** to cancel the model removal or select **Manage tables**.
Selecting **Manage tables** will take you to the **Tables** tab of the vector database console. Here you can choose to drop any tables that are preventing you from removing the model.

Manage Vector Tables Using the Vector Database Console

You can use the vector database console to perform management tasks, such as creating and dropping tables.

Once you are in the vector database console, select the **Tables** tab to view a list of existing tables. Each record in the list includes information about the table, such as the name, description, embedding model, and number of rows.

From the **Tables** page, you can choose to create new tables or manage existing ones.

- [Create a Vector Table](#)
- [Drop a Vector Table](#)

Create a Vector Table Using the Vector Database Console

Follow these steps to create a vector table using the vector database console.

1. From the vector database console **Overview** tab, select **Create table**.
2. Enter a **Table name**.
Specify the table name; it must consist of lowercase letters, numbers, and hyphens only.
3. Add a **Description** of the new table.
4. Select one of the following methods that you would like to use to supply data for the table.
 - **Loaded embedding model**
Select the model you would like to use from the drop down list of available loaded models.

If no models have been loaded, an error message will appear that prompts you to manage models. For more information about loading models, see [Load a Model Using the Vector Database Console](#).
 - **Bring my own vectors**
 - **Sample data**
Select from the available options which sample data you would like to use.
5. Select **Create**.

You will be directed to the **Tables** page that lists existing tables and displays their current status, such as **Active**, **Updating**, or **Failed**.

Drop a Vector Table Using the Vector Database Console

Follow these steps to drop a table using the vector database console.

To drop a table from your Vector Database, first navigate to the **Tables** tab of the vector database console.

1. Find the row of the table that you would like to drop and select the three dots  in the **ACTIONS** column.
2. Select **Drop table**.

Use the Query Playground

The query playground provides a user interface that makes it easy to test and explore vector similarity search.

With the query playground, you can perform text and vector searches on Vector Database tables. You can choose, using advanced options, a value for Top K Results to limit the query output, whether to include full vector data in the results, and to add filters to your search.

- [Perform a Text Search](#)
- [Perform a Vector Search](#)

Perform a Text Search

Perform a semantic search based on input query text using the query playground.

In order to perform a text search, you must have already loaded a model to your database. For instructions, see [Load a Model Using the Vector Database Console](#).

1. On the **Query playground** page, choose the table on which to perform the semantic search.
2. Select **Text search** and enter a query in the **Search query** text box.
This may be something like wireless headphones, which would produce results ordered by similarity to the search text.
3. Optionally enter a value for **Top K Results**. This will limit the number of records returned to the specified value.
4. Optionally click the checkbox to **Include vectors**.
If this option is selected, complete vector data will be included in the results.
5. Optionally specify **Filters**.
You may want to filter results on one or more fields. Enter the desired field, select an operation from the drop down list, and then provide a value.
Select **Add filter** to add filters or select the 'X' symbol to the right of the filter to remove it.
6. Click **Run search**.

Query results will be displayed by default in **Table view**. Click **JSON** to view the results in JSON format. In the **ACTIONS** column, click the three dots  to **Edit** or **Drop** a row. Selecting **Edit** takes you to a **Vector details** page, where you can view record information, edit metadata, and see associated vectors.

Perform a Vector Search

Perform a semantic search based on vector input using the query playground.

In order to perform a text search, you must have already loaded a model to your database. For instructions, see [Load a Model Using the Vector Database Console](#).

1. On the **Query playground** page, choose the table on which to perform the semantic search.
2. Select **Vector search** and enter vector value(s) in the text box. For example, you may input a vector in the format [0.1, 0.2, 0.3, ...].

3. Optionally enter a value for **Top K Results**. This will limit the number of records returned to the specified value.
4. Optionally click the checkbox to **Include vectors**.
If this option is selected, complete vector data will be included in the results.
5. Optionally specify **Filters**.
You may want to filter results on one or more fields. Enter the desired field, select an operation from the drop down list, and then provide a value.
Select **Add filter** to add filters or select the 'X' symbol to the right of the filter to remove it.
6. Click **Run search**.
Query results will be displayed by default in **Table view**. Click **JSON** to view the results in JSON format. In the **ACTIONS** column, click the three dots  to **Edit** or **Drop** a row. Selecting **Edit** takes you to a **Vector details** page, where you can view record information, edit metadata, and see associated vectors.

Manage Models

Use the included operations to manage models for use with your Vector Database.

- [List Available Models](#)
- [Describe a Model](#)
- [Load Your Own ONNX Model](#)
- [Drop a Model](#)

Load Your Own ONNX Model

Use the `load_model` operation to load an embedding or reranking model into the database.

The `load_model` operation imports a model from object storage, in ONNX format, for use in embedding generation and reranking operations. Once loaded, the model can be used for integrated table embeddings or for standalone inference.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

Caution

After loading an ONNX model using `DBMS_VECTOR_DATABASE`, drop the model only by using `DBMS_VECTOR_DATABASE.DROP_MODEL`.

Do not drop a model loaded through `DBMS_VECTOR_DATABASE` using `DBMS_DATA_MINING.DROP_MODEL`, `DBMS_VECTOR`, or another database API. Those APIs can remove the underlying database model without removing the `DBMS_VECTOR_DATABASE` metadata that references it.

This can leave an invalid model entry. For example, `DBMS_VECTOR_DATABASE.LIST_MODELS` might return a model whose details are `NULL`, and subsequent `DBMS_VECTOR_DATABASE` operations, such as `UPSERT_VECTORS`, can fail when they reference that model. To remove a model loaded through `DBMS_VECTOR_DATABASE`, use the following syntax:

```
BEGIN
  DBMS_VECTOR_DATABASE.DROP_MODEL('<model_name>');
END;
/
```

See the following for an example of using `load_model` to load a models from Oracle Object Storage:

- Load an embedding model:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>",
    password="<pass>"
))

response = client.load_model(
    model_name='all-MiniLM-L6-v2',
    url='https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace/b/bucket/o/model.onnx'
)
print(response)

#verify that the model was loaded using list_models
models = client.list_models()
print([item.model_name for item in models.items or []])
```

- Load a reranking model:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
```

```

        access_token="<bearer-token>", # or username="<user>",
        password="<pass>"
    ))

    response = client.load_model(
        model_name='reranker_model',
        url='https://objectstorage.example.com/models/reranker_quantized.onnx',
        model_params={'metadata': {'function': 'regression'}}
    )
    print(response)

    #verify that the model was loaded using list_models
    models = client.list_models()
    print([item.model_name for item in models.items or []])

```

A JSON response is returned confirming that the model was loaded successfully. An error is raised if the model already exists or if the provided URL is inaccessible.

Example response:

```

{
  "model_name": "SAMPLE_MODEL",
  "algorithm": "ONNX",
  "mining_function": "EMBEDDING",
  "creation_date": "2026-03-12T11:27:21Z",
  "attributes": [
    {
      "name": "DATA",
      "value": "TEXT",
      "data_type": "VARCHAR2",
      "data_length": 32767
    },
    {
      "name": "ORA$ONNXTARGET",
      "value": "VECTOR",
      "data_type": "VECTOR",
      "data_length": 1593,
      "vector_info": "VECTOR(384,FLOAT32)"
    }
  ]
}

```

For more information about the `load_model` operation, see [Python API Reference](#).

Curl

See how `POST /vecdb/models/` can be used in the following examples:

- Load an embedding model:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/models/" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

```

```
# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>" \

-d '{
  "modelName": "DOC_EMBED_MODEL",
  "url": "https://objectstorage.example.com/models/doc_embed.onnx"
}'
```

- Load a reranking model:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/models/" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>" \

-d '{
  "modelName": "RERANK_MODEL",
  "url": "https://objectstorage.example.com/models/
reranker_quantized.onnx",
  "modelParams": {
    "metadata": {
      "function": "regression"
    }
  }
}'
```

Responses:

- 201 Created – returns the registered model record with status information.

```
{
  "model_name": "DOC_EMBED_MODEL",
  "algorithm": "ONNX",
  "mining_function": "EMBEDDING",
  "creation_date": "2026-05-01T10:00:00.000000+00:00",
  "attributes": [
    {
      "name": "DATA",
      "value": "TEXT",
      "data_type": "VARCHAR2",
      "data_length": 4000,
      "vector_info": null
    }
  ]
}
```

- 400 - the request body included invalid parameters.
- 404 - model not found.

For more information about POST `/vecdb/models/`, see REST API Reference.

PL/SQL

See how `DBMS_VECTOR_DATABASE.LOAD_MODEL` can be used in the following examples.

- Load an embedding model:

```
dbms_vector_database.load_model(
    model_name => 'DOC_EMBED_MODEL',
    url => 'https://objectstorage.example.com/models/doc_embed.onnx'
);
```

- Load a reranking model:

```
dbms_vector_database.load_model(
    model_name => 'RERANK_MODEL',
    url => 'https://objectstorage.example.com/models/
reranker_quantized.onnx',
    model_params => JSON('{"metadata": {"function": "regression"}}')
);
```

Example response:

```
{
  "model_name": "DOC_EMBED_MODEL",
  "algorithm": "ONNX",
  "mining_function": "EMBEDDING",
  "creation_date": "2026-05-01T10:00:00.000000+00:00",
  "attributes": [
    {
      "name": "DATA",
      "value": "TEXT",
      "data_type": "VARCHAR2",
      "data_length": 4000,
      "vector_info": null
    }
  ]
}
```

For more information about the PL/SQL implementation, including parameters, see `LOAD_MODEL`.

List Available Models

Use the `list models` operation to list all loaded embedding and reranking models.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of using `list_models` to list loaded models and print their names:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_/db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

models = client.list_models()
print([item.model_name for item in models.items or []])
```

A JSON response is returned, containing an array of models, each with the following information:

- Model name
- Type (embedding, reranking)
- Algorithm
- Creation time stamp
- Attributes and parameters

Example response:

```
{
  "items": [
    {
      "model_name": "ALL_MINILM_L12_V2",
      "algorithm": "ONNX",
      "mining_function": "EMBEDDING",
      "creation_date": "2026-01-27T07:23:21Z",
      "attributes": [
        {
          "name": "DATA",
          "value": "TEXT",
          "data_type": "VARCHAR2",
          "data_length": 32767
        },
        {
          "name": "ORA$ONNXTARGET",
          "value": "VECTOR",
          "data_type": "VECTOR",
          "data_length": 1593,
          "vector_info": "VECTOR(384,FLOAT32)"
        }
      ]
    }
  ]
}
```

```
]
}
```

The string, `items`, is an array of model definitions returned by the service.

The `attributes` array includes model input and output attribute metadata. This includes `vector_info`, which is populated only when the value of `data_type` is `VECTOR`.

Pagination metadata is returned with `has_more`, `limit`, `offset`, and `count`.

For more information about the `list_models` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/models/` can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/models/" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Example 200 response:

```
{
  "items": [
    {
      "model_name": "DOC_EMBED_MODEL",
      "algorithm": "ONNX",
      "mining_function": "EMBEDDING",
      "creation_date": "2026-05-01T10:00:00.000000+00:00",
      "attributes": [
        {
          "name": "DATA",
          "value": "TEXT",
          "data_type": "VARCHAR2",
          "data_length": 4000,
          "vector_info": null
        }
      ]
    }
  ],
  "hasMore": false,
  "limit": 25,
  "offset": 0,
  "count": 1,
  "links": []
}
```

For more information about `GET /vecdb/models/`, see [REST API Reference](#).

PL/SQL

See how `DBMS_VECTOR_DATABASE.LIST_MODELS` can be used in the following example:

```
dbms_vector_database.list_models();
```

Example response:

```
{
  "models": [
    {
      "model_name": "DOC_EMBED_MODEL",
      "algorithm": "ONNX",
      "mining_function": "EMBEDDING",
      "creation_date": "2026-05-01T10:00:00.000000+00:00",
      "attributes": [
        {
          "name": "DATA",
          "value": "TEXT",
          "data_type": "VARCHAR2",
          "data_length": 4000,
          "vector_info": null
        }
      ]
    }
  ]
}
```

For more information about the PL/SQL implementation, including parameters, see `LIST_MODELS`.

Describe a Model

Use the `describe model` operation to retrieve detailed metadata for a loaded model.

The operation returns comprehensive information about the model, including its type, parameters, attributes, and usage statistics. You must include the name of the model as an input parameter.

An error is raised if the input model does not exist.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of using `describe_model`:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

details = client.describe_model(model_name='all-MiniLM-L12-v2')
print(details)
```

A JSON response is returned containing the following information:

- Model name and type
- Algorithm and mining function
- Input/output attributes
- Creation time stamp
- Model parameters

Example response:

```
{
  "model_name": "ALL_MINILM_L12_V2",
  "algorithm": "ONNX",
  "mining_function": "EMBEDDING",
  "creation_date": "2026-01-27T07:23:21Z",
  "attributes": [
    {
      "name": "DATA",
      "value": "TEXT",
      "data_type": "VARCHAR2",
      "data_length": 32767
    },
    {
      "name": "ORA$ONNXTARGET",
      "value": "VECTOR",
      "data_type": "VECTOR",
      "data_length": 1593,
      "vector_info": "VECTOR(384, FLOAT32)"
    }
  ]
}
```

The attributes array includes model input and output attribute metadata. This includes `vector_info`, which is populated only when the value of `data_type` is `VECTOR`.

For more information about the `describe_model` operation, see [Python API Reference](#).

Curl

See how GET /vecdb/models/{model_name} can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/models/DOC_EMBED_MODEL" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Responses:

- Example 200 response:

```
{
  "model_name": "DOC_EMBED_MODEL",
  "algorithm": "ONNX",
  "mining_function": "EMBEDDING",
  "creation_date": "2026-05-01T10:00:00.000000+00:00",
  "attributes": [
    {
      "name": "DATA",
      "value": "TEXT",
      "data_type": "VARCHAR2",
      "data_length": 4000,
      "vector_info": null
    }
  ]
}
```

- 404 - model not found.

For more information about GET /vecdb/models/{model_name}, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.DESCRIBE_MODEL can be used in the following example:

```
dbms_vector_database.describe_model('DOC_EMBED_MODEL');
```

Example response:

```
{
  "model_name": "DOC_EMBED_MODEL",
  "algorithm": "ONNX",
  "mining_function": "EMBEDDING",
  "creation_date": "2026-05-01T10:00:00.000000+00:00",
  "attributes": [
    {
      "name": "DATA",
```

```

        "value": "TEXT",
        "data_type": "VARCHAR2",
        "data_length": 4000,
        "vector_info": null
    }
]
}

```

For more information about the PL/SQL implementation, including parameters, see `DESCRIBE_MODEL`.

Drop a Model

Use the `drop_model` operation to remove a loaded model from the database.

The operation drops the specified embedding or reranking model. Models currently in use by vector tables cannot be dropped. An attempt to do so will raise an error.

You must provide the model name of the model to be dropped.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of using `drop_model`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

drop_mod_response = client.drop_model(model_name='sample_model')
print(drop_mod_response)

```

If the removal is successful, a JSON response is returned confirming the model deletion:

```

{
  "dropped": "SAMPLE_MODEL",
  "message": "Model SAMPLE_MODEL dropped successfully"
}

```

For more information about the `drop_model` operation, see [Python API Reference](#).

Curl

See how DELETE /vecdb/models/{model_name} can be used in the following example:

```
curl -X DELETE \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/models/
DOC_EMBED_MODEL" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Responses:

- 200 Created – returns a confirmation that the drop was successful:

```
{
  "message": "Model DOC_EMBED_MODEL dropped successfully"
}
```

- 400 - the model is used in a vector table and cannot be deleted.
- 404 - model not found.

For more information about DELETE /vecdb/models/{model_name}, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.DROP_MODEL can be used in the following example:

```
dbms_vector_database.drop_model('DOC_EMBED_MODEL');
```

Example response:

```
{
  "message": "Model DOC_EMBED_MODEL dropped successfully"
}
```

For more information about the PL/SQL implementation, including parameters, see DROP_MODEL.

Manage Vector Tables

Use the included operations to create and manage vector tables in your Vector Database.

- [Describe the Vector Database](#)
- [List Vector Tables](#)
- [Create a Vector Table](#)

- [Describe a Vector Table](#)
- [Drop a Vector Table](#)
- [Update a Vector Table Annotation](#)

Describe the Vector Database

Use the `describe_vector_database` operation to get summary statistics for the entire Vector Database service.

The operation returns a JSON response with database-level statistics, including total tables, models, and vectors.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of the syntax used to call `describe_vector_database`:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))
summary = client.describe_vector_database()
print(summary)
```

Example response:

```
{
  "total_tables": 48,
  "total_vectors": 1353894,
  "total_models": 9
}
```

For more information about the `describe_vector_database` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/summary` can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/summary" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
```

```
-u "<user>:<password>"

# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>"
```

Example 200 response:

```
{
  "total_tables": 3,
  "total_vectors": 12500,
  "total_models": 2
}
```

For more information about GET /vecdb/summary, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.SUMMARY() can be used in the following example:

```
dbms_vector_database.summary()
```

Example response:

```
{
  "total_tables": 3,
  "total_vectors": 12500,
  "total_models": 2
}
```

For more information about the PL/SQL implementation, including parameters, see SUMMARY.

List Vector Tables

Use the `list_vector_tables` operation to list all vector tables in the database.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See how `list_vector_tables` can be used in the following example:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
```

```

        rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
        access_token="<bearer-token>", # or username="<user>", password="<pass>"
    ))

tables = client.list_vector_tables()
print(tables)

```

A JSON response is returned, containing an array of table information accessible to the current user, including the following:

- Table names
- Descriptions
- Vector types
- Row counts
- Creation time stamps

For more information about the `list_vector_tables` operation, see [Python API Reference](#).

Curl

See how GET `/vecdb/vector-tables` can be used in the following example:

```

curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"

```

Example 200 response:

```

{
  "count": 1,
  "hasMore": false,
  "items": [
    {
      "table_name": "PRODUCT_VECTORS",
      "comment": "Product catalog embeddings",
      "table_params": {
        "auto_generate_id": false
      },
      "annotations": {
        "domain": "retail"
      },
      "vector_type": "dense",
      "vector_table_type": "BYOV",
      "embed_params": null,
      "index_params": {
        "vector_index_params": {
          "auto_index": true,

```

```

        "organization": "PARTITIONS",
        "distance_metric": "COSINE"
    },
    "metadata_index_params": {
        "auto_index": true
    },
    "parallel_creation": 1
},
"owner": "APPUSER",
"indexes": [],
"status": "Empty",
"stats": {
    "total_vectors": 0
},
"created": "2026-05-01T10:00:00.000000+00:00",
"updated": "2026-05-01T10:00:00.000000+00:00"
}
],
"limit": 25,
"offset": 0
}

```

For more information about GET /vecdb/vector-tables, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.LIST_VECTOR_TABLES() can be used in the following example:

```
dbms_vector_database.list_vector_tables();
```

Example response:

```

{
  "vector_tables": [
    {
      "table_name": "MYVECTORS",
      "comment": "Product embeddings",
      "table_params": {
        "auto_generate_id": false
      },
      "annotations": {
        "domain": "retail"
      },
      "vector_type": "dense",
      "vector_table_type": "BYOV",
      "embed_params": null,
      "index_params": {
        "vector_index_params": {
          "auto_index": true,
          "organization": "PARTITIONS",
          "distance_metric": "COSINE"
        },
        "metadata_index_params": {
          "auto_index": true
        }
      }
    }
  ]
}

```

```

        },
        "parallel_creation": 1
    },
    "owner": "APPUSER",
    "indexes": [],
    "status": "Empty",
    "stats": {
        "total_vectors": 0
    },
    "created": "2026-05-01T10:00:00.000000+00:00",
    "updated": "2026-05-01T10:00:00.000000+00:00"
}
]
}

```

For more information about the PL/SQL implementation, including parameters, see `LIST_VECTOR_TABLES`.

Create a Vector Table

Use the `create vector table` operation to create a new vector table for storing vector embeddings.

The operation creates a vector table with a fixed schema optimized for vector search. The table includes columns for ID, vector data, and JSON metadata. You can configure automatic ID generation, embedding integration, and index parameters during creation.

If the table already exists or if invalid parameters are provided, an error is raised.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following example for different variations of table creation using `create_vector_table()`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#create a table for pre-computed vectors
prod_vec_response = client.create_vector_table(
    name="product_vectors",
    comment="Product embeddings",
    table_params={"auto_generate_id": True},

```

```

        index_params={
            "vector_index_params": {
                "auto_index": True,
                "organization": "PARTITIONS",
                "distance_metric": "COSINE",
            }
        },
    )
print(prod_vec_response)

#create a table with integrated embedding
docs_response = client.create_vector_table(
    name="documents",
    table_params={"auto_generate_id": True},
    embed_params={
        "model": "all_MiniLM_L12_v2",
        "embed_metadata_jsonpath": "content",
    },
)
print(docs_response)
#create a table for bring-your-own vectors with cust_vec_response =
client.create_vector_table(
    cust_vec_response = client.create_vector_table(
        name="customer_vectors",
        comment="Manually managed vector table",
        index_params={
            "vector_index_params": {
                "auto_index": False,
            }
        },
    )
)
print(cust_vec_response)

```

A JSON response containing the created table details and status is returned.

The following parameters are accepted as input:

Parameter	Type	Description
table_name	str	Name of the vector table to create. The provided name must be unique within the database.
description	str	An optional, human-readable description of the table's purpose.
auto_generate_id	bool	If set to true, automatically generates unique IDs for vectors on insert. If set to false, IDs must be provided explicitly. This is an optional parameter and is set by default to false.
annotations	str	You can optionally provide key-value pairs for custom metadata about the table. For example: <pre> { 'application': 'chatbot', 'department': 'sales' } </pre>

Parameter	Type	Description
vector_type	str	Optionally specify the type of vectors to store. Currently, only dense vectors are supported. The default value is dense.
embed_params	dict	Optionally provide configuration for an integrated embedding mode. If provided, the table will automatically generate embeddings on insert. For example: <pre>{ 'model': 'all_MiniLM_L12_v2', 'embed_metadata_jsonpath': 'content' }</pre>
index_params	dict	Optionally provide configuration for vector index creation. This parameter supports both automatic and manual indexing strategies. For example: For auto-managed IVF indexes: <pre>{ 'indexing': 'auto', 'organization': 'PARTITIONS', 'advanced_params': { 'type': 'IVF', 'partitions': 5 } }</pre> To defer creation and later supply detailed settings (such as distance metric, organization, IVF/HNSW parameters) when using the create_index operation: <pre>{ 'indexing': 'manual' }</pre>
debug_flags	dict	Optionally provide debug configuration for detailed logging.

Curl

See the following examples of table creation using POST /vecdb/vector-tables:

- Create a table for bring-your-own vectors:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
tables/" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \
```

```
# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>" \

-d '{
  "name": "product_vectors",
  "comment": "Product catalog embeddings",
  "tableParams": {
    "auto_generate_id": false
  },
  "annotations": {
    "domain": "retail"
  },
  "indexParams": {
    "vector_index_params": {
      "auto_index": true,
      "organization": "PARTITIONS",
      "distance_metric": "COSINE",
      "accuracy": 90
    },
    "metadata_index_params": {
      "auto_index": true,
      "include_paths": ["category", "price"]
    },
    "parallel_creation": 2
  }
}'
```

- Create a model-backed table:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
tables/" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>" \

-d '{
  "name": "product_text_vectors",
  "comment": "Product text embeddings",
  "tableParams": {
    "auto_generate_id": true
  },
  "embedParams": {
    "model": "DOC_EMBED_MODEL",
    "embed_metadata_jsonpath": "description"
  },
  "indexParams": {
    "vector_index_params": {
      "auto_index": true,
      "organization": "INMEMORY GRAPH",
```

```

        "distance_metric": "COSINE"
    }
}
}'

```

Responses:

- Example 201 response:

```

{
  "table_name": "PRODUCT_VECTORS",
  "comment": "Product catalog embeddings",
  "table_params": {
    "auto_generate_id": false
  },
  "annotations": {
    "domain": "retail"
  },
  "vector_type": "dense",
  "vector_table_type": "BYOV",
  "embed_params": null,
  "index_params": {
    "vector_index_params": {
      "auto_index": true,
      "organization": "PARTITIONS",
      "distance_metric": "COSINE",
      "accuracy": 90
    },
    "metadata_index_params": {
      "auto_index": true,
      "include_paths": ["category", "price"]
    },
    "parallel_creation": 2
  },
  "owner": "APPUSER",
  "indexes": [],
  "status": "Empty",
  "stats": {
    "total_vectors": 0
  },
  "created": "2026-05-01T10:00:00.000000+00:00",
  "updated": "2026-05-01T10:00:00.000000+00:00"
}

```

- 400 - invalid or missing parameters, including invalid nested object fields or unsupported combinations.
- 409 - a table with that name already exists.

The following parameters are accepted as part of the request body:

Parameter	Type	Description
name	string	Name of the vector table to create. The provided name must be unique within the database.
comment	string	An optional, human-readable description of the table's purpose.
annotations	object	You can optionally provide key-value pairs for custom metadata about the table.

Parameter	Type	Description
tableParams	object	Optionally provide table-level creation parameters, such as <code>auto_generate_id</code> . If set to <code>true</code> , <code>auto_generate_id</code> automatically generates unique IDs for vectors on insert. If set to <code>false</code> , IDs must be provided explicitly. This is an optional parameter and is set by default to <code>false</code> .
embedParams	object	Optionally specify automatic embedding configuration for model-backed tables. Supply this parameter when the table should generate vectors from metadata text by using a loaded embedding model. Omit for bring-your-own-vector tables.
indexParams	object	Optionally set vector and metadata index configuration. The following fields can be specified: <code>index_type</code> <code>vector_index_params</code> <code>metadata_index_params</code> <code>parallel_creation</code>
debugFlags	object	Optionally provide debug configuration for detailed logging.

Response fields:

Field	Type	Description
table_name	string	The name of the table.
comment	string	Description of the vector table. Nullable.
auto_generate_id	integer	Whether the ID column is auto-generated.
owner	string	Database owner of the table.
status	string	Population status of the vector table.
vector_table	string	"dense"
vector_table_type	string	"BYOV" or "MODEL"
index_params	object	Index configuration with fields <code>indexing</code> , <code>organization</code> , <code>distance_metric</code> , <code>accuracy</code> , and <code>advanced_params</code> . Nullable.
annotations	object or string	User metadata attached to the table. Nullable.
embed_params	object	Embedding model configuration for MODEL tables, with fields <code>model</code> and <code>embed_metadata_jsonpath</code> . Nullable.
dense_idx_name	string	Name of the vector index associated to the table. Nullable.
stats	object	Table statistics. Contains <code>total_vectors</code> (number, nullable).
created	string (date-time)	Creation timestamp.
updated	string (date-time)	Last updated timestamp. Nullable.

For more information, see [Create a vector table](#).

PL/SQL

See how DBMS_VECTOR_DATABASE.CREATE_VECTOR_TABLE can be used in the following example:

```
dbms_vector_database.create_vector_table(  
    name => 'product_vectors',  
    comment => 'Product catalog embeddings',  
    table_params => JSON('{"auto_generate_id": false}'),  
    annotations => JSON('{"domain": "retail"}'),  
    index_params => JSON('{  
        "vector_index_params": {  
            "auto_index": true,  
            "organization": "PARTITIONS",  
            "distance_metric": "COSINE",  
            "accuracy": 90  
        },  
        "metadata_index_params": {  
            "auto_index": true,  
            "include_paths": ["category", "price"]  
        },  
        "parallel_creation": 2  
    }')  
);
```

Example response:

```
{  
    "table_name": "PRODUCT_VECTORS",  
    "comment": "Product catalog embeddings",  
    "table_params": {  
        "auto_generate_id": false  
    },  
    "annotations": {  
        "domain": "retail"  
    },  
    "vector_type": "dense",  
    "vector_table_type": "BYOV",  
    "embed_params": null,  
    "index_params": {  
        "vector_index_params": {  
            "auto_index": true,  
            "organization": "PARTITIONS",  
            "distance_metric": "COSINE",  
            "accuracy": 90  
        },  
        "metadata_index_params": {  
            "auto_index": true,  
            "include_paths": ["category", "price"]  
        },  
        "parallel_creation": 2  
    },  
    "owner": "APPUSER",  
    "indexes": [],  
    "status": "Empty",  
    "stats": {  
        "total_vectors": 0  
    }  
}
```

```
    },  
    "created": "2026-05-01T10:00:00.000000+00:00",  
    "updated": "2026-05-01T10:00:00.000000+00:00"  
}
```

Describe a Vector Table

Use the `describe_vector_table` operation to retrieve detailed configuration and metadata for a vector table.

The operation returns comprehensive information about the specified table, including its schema, index configuration, embedding settings, row count, and creation time stamp.

A table name is required as input. An error is raised if the provided table does not exist.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of using `describe_vector_table` to request table information:

```
from oracle_vecdb import OracleVecDB, Configuration  
  
client = OracleVecDB(Configuration(  
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",  
    access_token="<bearer-token>", # or username="<user>", password="<pass>"  
))  
  
details = client.describe_vector_table(name="product_vectors")  
print(details.to_dict())
```

A JSON response containing table configuration is returned, including the following information:

- Table name and description
- Vector type and dimensions
- Index parameters and status
- Embedding model configuration (if applicable)
- Row count and storage statistics
- Annotations and metadata

Example response:

```
{
  "table_name": "DEMO_PRODUCTS",
  "owner": "VECTOR3",
  "status": "Empty",
  "vector_type": "dense",
  "vector_table_type": "BYOV",
  "index_params": {
    "vector_index_params": {
      "auto_index": true,
      "organization": "PARTITIONS",
      "advanced_params": {
        "partitions": 10
      }
    },
    "parallel_creation": 4
  },
  "annotations": {
    "metric": "cosine",
    "dimension": "5"
  },
  "embed_params": null,
  "created": "2026-03-12T11:21:30Z",
  "description": "Demo table for SDK examples"
}
```

For more information about the `describe_vector_table` operation, see [Python API Reference](#).

Curl

See the following for an example of using `GET /vecdb/vector-tables/{vector_table_name}` to request table information:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/
  product_vectors" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Responses:

- Example 200 response:

```
{
  "table_name": "PRODUCT_VECTORS",
  "comment": "Product catalog embeddings",
  "table_params": {
    "auto_generate_id": false
  }
}
```

```

    },
    "annotations": {
      "domain": "retail"
    },
    "vector_type": "dense",
    "vector_table_type": "BYOV",
    "embed_params": null,
    "index_params": {
      "vector_index_params": {
        "auto_index": true,
        "organization": "PARTITIONS",
        "distance_metric": "COSINE"
      },
      "metadata_index_params": {
        "auto_index": true
      }
    },
    "parallel_creation": 1
  },
  "owner": "APPUSER",
  "indexes": [],
  "status": "Empty",
  "stats": {
    "total_vectors": 0
  },
  "created": "2026-05-01T10:00:00.000000+00:00",
  "updated": "2026-05-01T10:00:00.000000+00:00"
}

```

- 404 - not found or precondition failed.

For more information about GET /vecdb/vector-tables/{vector_table_name}, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.DESCRIBE_VECTOR_TABLE can be used in the following example:

```
dbms_vector_database.describe_vector_table('product_vectors');
```

Example response:

```

{
  "table_name": "PRODUCT_VECTORS",
  "comment": "Product catalog embeddings",
  "table_params": {
    "auto_generate_id": false
  },
  "annotations": {
    "domain": "retail"
  },
  "vector_type": "dense",
  "vector_table_type": "BYOV",
  "embed_params": null,
  "index_params": {
    "vector_index_params": {

```

```

        "auto_index": true,
        "organization": "PARTITIONS",
        "distance_metric": "COSINE"
    },
    "metadata_index_params": {
        "auto_index": true
    },
    "parallel_creation": 1
},
"owner": "APPUSER",
"indexes": [],
"status": "Empty",
"stats": {
    "total_vectors": 0
},
"created": "2026-05-01T10:00:00.000000+00:00",
"updated": "2026-05-01T10:00:00.000000+00:00"
}

```

For more information about the PL/SQL implementation, including parameters, see `DESCRIBE_VECTOR_TABLE`.

Drop a Vector Table

Use the `drop_vector_table` operation to permanently delete a vector table and all of its data. The operation drops the specified vector table, including all vectors, metadata, and associated indexes.

Caution

This operation is irreversible. All data in the table will be permanently deleted.

A table name is required as input. An error is raised if the provided table does not exist.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of dropping a vector table using `drop_vector_table`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_/_/db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"

```

```

))

response = client.drop_vector_table(name="old_vectors")
print(response)

```

A JSON response is returned, confirming the table was dropped successfully.

Example response:

```

{
  "message": "Table DEMO_PRODUCTS deleted successfully"
}

```

For more information about the `drop_vector_table` operation, see [Python API Reference](#).

Curl

See the following example of dropping a table using `DELETE /vecdb/vector-tables/{vector_table_name}`:

```

curl -X DELETE \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/product_vectors" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"

```

Responses:

- Example 200 response:

```

{
  "message": "Table PRODUCT_VECTORS deleted successfully"
}

```

- 404 - table not found.

For more information about `DELETE /vecdb/vector-tables/{vector_table_name}`, see [REST API Reference](#).

PL/SQL

See how `DBMS_VECTOR_DATABASE.DROP_MODEL` can be used in the following example:

```

dbms_vector_database.drop_vector_table('product_vectors');

```

Example response:

```
{
  "message": "Table PRODUCT_VECTORS deleted successfully"
}
```

For more information about the PL/SQL implementation, including parameters, see `DROP_VECTOR_TABLE`.

Update a Vector Table Annotation

Use the `update_vector_table_annotation` operation to update the description and annotations for an existing vector table.

The operation modifies the metadata and configuration of a vector table without affecting the stored data. You can update the description, annotations, and index parameters.

An error is raised if the table does not exist or if invalid parameters are provided.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following example of updating a table's annotations using `update_vector_table_annotation`:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

response = client.update_vector_table_annotation(
    name='products',
    comment='Updated product embeddings',
    annotations={'version': '2.0', 'updated': '2025-01-27'}
)
print(response)
```

A JSON response is returned confirming the update.

Example response:

```
{
  "table_name": "products",
  "annotations": {
```

```

        "version": "2.0",
        "updated": "2025-01-27"
    },
    "description": "Product embeddings"
}

```

For more information about the `update_vector_table_annotation` operation, see [Python API Reference](#).

Curl

See how `PATCH /vecdb/vector-tables/{vector_table_name}` can be used in the following example:

```

curl -X PATCH \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/product_vectors" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "comment": "Updated product embeddings",
    "annotations": {
      "domain": "retail",
      "owner": "search"
    }
  }'

```

Responses:

- Example 202 response:

```

{
  "table_name": "PRODUCT_VECTORS",
  "comment": "Updated product embeddings",
  "table_params": {
    "auto_generate_id": false
  },
  "annotations": {
    "domain": "retail",
    "owner": "search"
  },
  "vector_type": "dense",
  "vector_table_type": "BYOV",
  "embed_params": null,
  "index_params": {
    "vector_index_params": {
      "auto_index": true
    }
  }
}

```

```

    },
    "metadata_index_params": {
      "auto_index": true
    },
    "parallel_creation": 1
  },
  "owner": "APPUSER",
  "indexes": [],
  "status": "Empty",
  "stats": {
    "total_vectors": 0
  },
  "created": "2026-05-01T10:00:00.000000+00:00",
  "updated": "2026-05-02T09:30:00.000000+00:00"
}

```

- 400 - invalid JSON, invalid types, or missing required fields.
- 404 - not found or precondition failed.

For more information about PATCH `/vecdb/vector-tables/{vector_table_name}`, see REST API Reference.

PL/SQL

See how `DBMS_VECTOR_DATABASE.UPDATE_VECTOR_TABLE_ANNOTATION` can be used in the following example:

```

dbms_vector_database.update_vector_table_annotation(
  name => 'product_vectors',
  comment => 'Updated product embeddings',
  annotations => JSON('{"domain": "retail", "owner": "search"}')
);

```

Example response:

```

{
  "table_name": "PRODUCT_VECTORS",
  "comment": "Updated product embeddings",
  "table_params": {
    "auto_generate_id": false
  },
  "annotations": {
    "domain": "retail",
    "owner": "search"
  },
  "vector_type": "dense",
  "vector_table_type": "BYOV",
  "embed_params": null,
  "index_params": {
    "vector_index_params": {
      "auto_index": true
    },
    "metadata_index_params": {
      "auto_index": true
    }
  },
}

```

```

    "parallel_creation": 1
  },
  "owner": "APPUSER",
  "indexes": [],
  "status": "Empty",
  "stats": {
    "total_vectors": 0
  },
  "created": "2026-05-01T10:00:00.000000+00:00",
  "updated": "2026-05-02T09:30:00.000000+00:00"
}

```

For more information about the PL/SQL implementation, including parameters, see [UPDATE_VECTOR_TABLE_ANNOTATION](#).

Manage Data

Upsert records, load data sets, and perform other vector management tasks in your Vector Database using the included operations.

- [Embed Data](#)
- [Upsert Records](#)
- [Load Data Sets](#)
- [List Vectors](#)
- [Delete Vectors](#)
- [List Vector Load Jobs](#)
- [Describe a Bulk Load Job](#)
- [Retrieve Log Output for Bulk Loads](#)

Embed Data

Use the `generate_embedding` operation to generate embeddings (vector representation) for provided inputs using a specified embedding model.

This operation converts text into dense vector representations using the specified embedding model. Ensure to load the model in the database using [Load Your Own ONNX Model](#) before initiating `generate_embedding` operation. Loading the model is a one-time operation; you do not need to reload it every time while using `generate_embedding` operation to generate embeddings.

Note

This operation runs the loaded embedding model within Oracle AI Database and consumes database CPU resources. Large inputs, large batches, or frequent calls can affect database workload performance. For embedding-heavy workloads, consider using OCI Generative AI to generate embeddings externally.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of how to use the `generate_embedding` operation to generate embeddings.

Begin by loading a vector embedding model into your database, then generate the embeddings and inspect the response.

```
from oracle_vecdb import OracleVecDB, Configuration

client.load_model(
    model_name="all-MiniLM-L6-v2",
    url="https://objectstorage.example.com/bucket/all-MiniLM-L6-v2.onnx",
)

embeddings = client.generate_embedding(
    model_name="all-MiniLM-L6-v2",
    inputs=[
        "Wireless noise-cancelling headphones",
        "Ergonomic office chair with lumbar support",
    ],
)

print(embeddings.to_dict()["data"][0]["embedding"])
```

A JSON response is returned containing generated embeddings, an array of vectors corresponding to each input. An error is raised if the model is not loaded or if inputs are invalid.

Example response:

```
{
  "data": [
    {
      "text": "Sample text to embed",
      "embedding": [
        0.0117027815,
        0.00419755466,
        -0.0301191173,
        0.011317092,
        0.0763486549
      ]
    }
  ]
}
```

For more information about the `generate_embedding` operation, see [Python API Reference](#).

Curl

See how POST /vecdb/embed can be used in the following example:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/embed" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "modelName": "DOC_EMBED_MODEL",
    "inputs": [
      {"text": "hello world"}
    ]
  }'
```

Responses:

- Example 200 response:

```
{
  "data": [
    {
      "text": "hello world",
      "embedding": [
        0.0117027815,
        0.00419755466,
        -0.0301191173,
        0.011317092
      ]
    }
  ]
}
```

- 400 - the request body included invalid parameters.

For more information about POST /vecdb/embed, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.GENERATE_EMBEDDING can be used in the following example:

```
dbms_vector_database.generate_embedding(
  model_name => 'DOC_EMBED_MODEL',
  inputs => SYS.JSON_ARRAY_T(['{"text": "hello world"}'])
);
```

Example output:

```
{
  "data": [
    {
      "text": "hello world",
      "embedding": [
        0.0117027815,
        0.00419755466,
        -0.0301191173,
        0.011317092
      ]
    }
  ]
}
```

For information about the PL/SQL implementation, including parameters, see `GENERATE_EMBEDDING`.

Upsert Records

Use the `upsert_vectors` operation to insert or update vectors in a table.

The operation upserts vectors into the specified table. If a vector with the same ID already exists, it will be updated with the new values. Otherwise, a new record is inserted. An error is raised if the provided table does not exist or if the vector format is invalid.

See the included sections for information about upserting text and vectors.

- [Upsert Text](#)
- [Upsert Vectors](#)

Upsert Text

For tables with embedding models, text can be upserted using the `metadata` field of the `vectors` parameter of the `upsert_vectors` operation.

For tables with integrated embedding models, you can provide text in the `metadata` field as one of the values, with `embed_metadata_jsonpath` configured upon table creation as the key. Vector embeddings will be generated automatically.

See [Create a Vector Table](#) for more information about using the `embed_params` parameter to configure the table for automatic embedding generation.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of upserting text into a table with automatic embedding using `upsert_vectors`. In this example, the table was created with the `embed_params` parameter configured with `"embed_metadata_jsonpath" : "content"`, so that any text provided in the `content` field will be auto-embedded.

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_/db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

ups_text_response = client.upsert_vectors(
    table_name='documents',
    vectors=[
        {
            'id': 'doc_1',
            'metadata': {
                'content': 'Machine learning is transforming healthcare',
                'category': 'AI',
                'author': 'John Doe'
            }
        },
        {
            'id': 'doc_2',
            'metadata': {
                'content': 'Vector databases enable semantic search',
                'category': 'Database',
                'author': 'Jane Smith'
            }
        }
    ]
)
print(ups_text_response)
```

If the table is configured with `auto_generate_id` set to `true`, then you don't need to provide `id` as part of the upsert object.

A JSON response is returned confirming successful upsert with a count of inserted and updated vectors.

Example response:

```
{
  "upserted_count": 10
}
```

For more information about the `upsert_vectors` operation, see [Python API Reference](#).

Curl

See the following example of upserting text using POST `/vecdb/vector-tables/{vector_table_name}/upsert`:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/
products/upsert" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>" \

-d '{
  "vectors": [
    {
      "id": "prod_1",
      "dense_vector": [0.1, 0.2, 0.3, 0.4, 0.5],
      "metadata": {
        "name": "Wireless Headphones",
        "category": "electronics",
        "price": 99.99
      }
    },
    {
      "id": "prod_2",
      "dense_vector": [0.2, 0.3, 0.1, 0.5, 0.4],
      "metadata": {
        "name": "Smart Watch",
        "category": "electronics",
        "price": 199.99
      }
    }
  ]
}'
```

Responses:

- Example 201 response:

```
{
  "upserted_count": 2
}
```

- 400 - the request body included invalid parameters.
- 404 - the specified vector table does not exist.

For more information about POST `/vecdb/vector-tables/{vector_table_name}/upsert`, see REST API Reference.

PL/SQL

For information about the PL/SQL implementation, including parameters, see `UPSERT_VECTORS`.

Upsert Vectors

Vectors can be inserted or updated in a table by listing them in the `vectors` parameter of the `upsert_vectors` operation.

You can include the dimension values of each vector using the `dense_vector` field of the `vectors` parameter.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of upserting pre-computed vectors:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

ups_vec_response = client.upsert_vectors(
    table_name='products',
    vectors=[
        {
            'id': 'prod_1',
            'dense_vector': [0.1, 0.2, 0.3, 0.4, 0.5],
            'metadata': {
                'name': 'Wireless Headphones',
                'category': 'electronics',
                'price': 99.99
            }
        },
        {
            'id': 'prod_2',
            'dense_vector': [0.2, 0.3, 0.1, 0.5, 0.4],
            'metadata': {
                'name': 'Smart Watch',
                'category': 'electronics',
                'price': 199.99
            }
        }
    ]
)
```

```
)  
print(ups_vec_response)
```

If the table is configured with `auto_generate_id` set to `true`, then you don't need to provide `id` as part of the upsert object.

A JSON response is returned confirming successful upsert with a count of inserted and updated vectors.

Example response:

```
{  
  "upserted_count": 10  
}
```

For more information about the `upsert_vectors` operation, see [Python API Reference](#).

Curl

See the following example of upserting vectors using `POST /vecdb/vector-tables/{vector_table_name}/upsert`:

```
curl -X POST \  
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/  
sample_table/upsert" \  
  -H "Content-Type: application/json" \  
  -H "Accept: application/json" \  
  # Choose ONE authentication method:  
  
  # Option 1: Basic authentication  
  -u "<user>:<password>" \  
  
  # Option 2: OAuth Bearer token  
  # -H "Authorization: Bearer <access_token>" \  
  
  -d '{  
    "vectors": [  
      {  
        "id": "vec1",  
        "dense_vector": [0.1, 0.5, 0.4, 0.56],  
        "metadata": {  
          "type": "document",  
          "source": "Wikipedia"  
        }  
      },  
      {  
        "id": "vec2",  
        "dense_vector": [0.14, 0.534, 0.12, 0.58],  
        "metadata": {  
          "type": "document",  
          "source": "Internal"  
        }  
      }  
    ]  
  }'
```

Responses:

- Example 201 response:

```
{
  "upserted_count": 2
}
```

- 400 - the request body included invalid parameters.
- 404 - the specified vector table does not exist.

For more information about POST /vecdb/vector-tables/{vector_table_name}/upsert, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.UPSERT_VECTORS can be used in the following example:

```
dbms_vector_database.upsert_vectors(
  'sample_table',
  vectors => JSON('[
    {
      "id": "vec1",
      "dense_vector": [0.1, 0.5, 0.4, 0.56],
      "metadata": {"type": "document", "source": "Wikipedia"}
    },
    {
      "id": "vec2",
      "dense_vector": [0.14, 0.534, 0.12, 0.58],
      "metadata": {"type": "document", "source": "Internal"}
    }
  ]')
);
```

Example response:

```
{
  "upserted_count": 2
}
```

For more information and parameters, see UPSERT_VECTORS.

Load Data Sets

Use the load_vectors operation to perform bulk loads of vectors from a CSV file in Oracle Object Storage.

The operation appends vector records from Oracle Object Storage into an existing vector table asynchronously. The table must already exist. For integrated embedding tables, the table's embed_params configuration is used to generate embeddings when needed.

The object storage URL provided as a parameter should point to a CSV file with the following format:

```
id,dense_vector,metadata
id1,[0.1, 0.2, 0.3],{"field1": "value1", "field2": "value2"}
id2,[0.4, 0.5, 0.6],{"field1": "value3", "field2": "value4"}
```

Guidelines for preparing CSV files for loading data set:

- General Structure
 - The file should be a valid, comma-separated values (CSV) file.
 - Every file should begin with a header row that defines column names.
 - Each field should be separated by a comma (,).
- Handling Fields with Commas
 - If a field (such as a dense vector or metadata) contains commas, enclose the entire field in double quotes ("). This ensures the CSV parser treats the entire value as a single column.
- Embedding JSON in CSV: If the following guidelines are not followed, the CSV parser may incorrectly split fields, resulting in ingestion errors.
 - JSON fields should use double quotes (") for all property names and string values.
 - To include double quotes inside JSON within a CSV field, escape them by doubling: (").
 - Do not use single quotes (') in JSON.
 - Always enclose JSON content in double quotes (").
 - JSON example:

```
id,metadata 77E0D7F0-1942-494A-
ACE2-9004D2BDC59E,"{"PARK_CODE":"abli","NAME":"Abraham Lincoln
Birthplace","STATES":"KY"}"
```
- Header Row Formats
 - For embedded model table:
 - * ID, METADATA)
 - * METADATA (if auto_generated_id is set to True
 - For bring your own vector table:
 - * ID, DENSE_VECTOR
 - * ID,DENSE_VECTOR, METADATA
 - * METADATA (if auto_generated_id is set to True)
 - * DENSE_VECTOR (if auto_generated_id is set to True)
 - The order of columns is flexible.
 - Headers are not case-sensitive (for example: METADATA, Metadata, or "Metadata" are all valid).
- ID Field Guidelines
 - If auto_generated_id is set to False, each data row should include a unique ID in the appropriate column.

- If `auto_generated_id` is set to `True`, providing the ID column is optional.

A JSON response is returned containing the load job ID and initial status.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of bulk loading vectors from Oracle Object Storage and monitoring the load process.

First load vectors from object storage into a variable called `response`, then use `describe_vector_load_job` to monitor the load process:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

response = client.load_vectors(
    table_name='products',
    url='https://objectstorage.region.oraclecloud.com/.../vectors.csv',
    params={'credential': 'OCI_CREDENTIAL'}
)

import json
job = json.loads(response)
status = client.describe_vector_load_job(job['job_name'])

print(status)
```

Example response:

```
{
  "job_name": "VECDB_LOAD_ABC123",
  "job_creator": "VECTOR3",
  "operation": "LOAD",
  "state": "SUCCEEDED",
  "links": [
    {
      "rel": "collection",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/load/jobs/"
    },
    {
      "rel": "self",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/load/jobs/vecdb_load_abc123/"
    }
  ]
}
```

```

    {
      "rel": "related",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/load/jobs/vecdb_load_abc123/jobfile"
    }
  ]
}

```

For more information about the `load_vector`'s operation, see [Python API Reference](#).

Curl

See how `POST /vecdb/load` can be used in the following example:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/load" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "tableName": "product_vectors",
    "url": "https://objectstorage.example.com/data/product_vectors.csv",
    "params": {
      "credential": "OBJ_STORE_CRED"
    }
  }'

```

Responses:

- Example 200 response:

```

{
  "job_creator": "APPUSER",
  "job_name": "VECDB_LOAD_20260501110000",
  "job_type": "PLSQL_BLOCK",
  "operation": "LOAD",
  "state": "SCHEDULED",
  "start_date": "2026-05-01T11:00:00.118886Z",
  "links": [
    {
      "href": "/vecdb/load/jobs/",
      "rel": "collection"
    },
    {
      "href": "/vecdb/load/jobs/vecdb_load_20260501110000/",
      "rel": "self"
    }
  ]
}

```

```

        "href": "/vecdb/load/jobs/vecdb_load_20260501110000/jobfile",
        "rel": "related"
    }
  ]
}

```

- 400 - the request body included invalid parameters.

For more information about POST /vecdb/load, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.LOAD_VECTORS can be used in the following example:

```

dbms_vector_database.load_vectors(
    table_name => 'product_vectors',
    url => 'https://objectstorage.example.com/data/product_vectors.csv',
    params => JSON('{"credential": "OBJ_STORE_CRED"}')
);

```

Example response:

```

{
  "table_name": "PRODUCT_VECTORS",
  "file_url": "https://objectstorage.example.com/data/product_vectors.csv"
}

```

For more information about the PL/SQL implementation, including parameters, see LOAD_VECTORS.

List Vectors

Use the `list_vectors` operation to retrieve vectors from a table by their IDs.

The operation lists vector records with their IDs, embeddings, and metadata. Pagination is supported for large result sets.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for examples of how to implement `list_vectors`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",

```

```

        access_token="<bearer-token>", # or username="<user>", password="<pass>"
    ))

    #get specific vectors by ID
    vectors = client.list_vectors(
        table_name='products',
        ids=['prod_1', 'prod_2', 'prod_3']
    )
    print(vectors)

    #paginate through results, including only the first 10 results
    page1 = client.list_vectors(
        table_name='products',
        limit=10,
        offset=0
    )
    print(page1)

    #then, paginate through the next 10 results
    page2 = client.list_vectors(
        table_name='products',
        limit=10,
        offset=10
    )
    print(page2)

```

A JSON response is returned containing matching vectors with the following information:

- IDs
- Dense vectors
- Metadata

Example response:

```

{
  "items": [
    {
      "id": "prod_010",
      "dense_vector": [0.21, 0.19, 0.24, 0.47, 0.40],
      "metadata": {
        "name": "Halo Headphones",
        "category": "electronics",
        "price": 199.0,
        "color": "onyx"
      }
    },
    {
      "id": "prod_003",
      "dense_vector": [0.08, 0.25, 0.34, 0.41, 0.50],
      "metadata": {
        "name": "Zephyr Running Tee",
        "category": "apparel",
        "price": 39.95,
        "color": "light grey"
      }
    }
  ]
}

```

```

    }
  ],
  "limit": 15,
  "offset": 0,
  "count": 10
}

```

Curl

See how POST /vecdb/vector-tables/{vector_table_name}/list can be used in the following examples:

- Get specific vectors by ID:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/sample_table/list" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  -H "Authorization: Bearer <access_token>" \

  -d '{
    "ids": ["id1", "id2"],
    "limit": 2,
    "offset": 0
  }'

```

- Paginate through results (first 10 results):

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/sample_table/list" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  -u "<user>:<password>" \
  -d '{
    "limit": 10,
    "offset": 0
  }'

```

- Paginate through the next 10 results:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/sample_table/list" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  -u "<user>:<password>" \
  -d '{
    "limit": 10,

```

```
    "offset": 10
  }'
```

Responses:

- Example 200 response:

```
{
  "items": [
    {
      "id": "id1",
      "dense_vector": [0.1, 0.5, 0.4, 0.56],
      "metadata": {
        "product_name": "Noise Cancelling Headphones",
        "category": "electronics",
        "brand": "Sony",
        "price": 299.99
      }
    },
    {
      "id": "id2",
      "dense_vector": [0.14, 0.534, 0.12, 0.58],
      "metadata": {
        "product_name": "Studio Monitor Headphones",
        "category": "electronics",
        "brand": "Audio-Technica",
        "price": 149.99
      }
    }
  ],
  "limit": 2,
  "offset": 0,
  "count": 2
}
```

- 400 - the request body included invalid parameters.
- 404 - the specified vector table does not exist.

PL/SQL

See how DBMS_VECTOR_DATABASE.LIST_VECTORS can be used in the following example:

```
dbms_vector_database.list_vectors(
  table_name => 'sample_table',
  ids => JSON('["id1", "id2"]')
);
```

Example response:

```
{
  "items": [
    {
      "id": "id1",
      "dense_vector": [0.1, 0.5, 0.4, 0.56],
      "metadata": {
```

```

        "product_name": "Noise Cancelling Headphones",
        "category": "electronics",
        "brand": "Sony",
        "price": 299.99
    }
},
{
    "id": "id2",
    "dense_vector": [0.14, 0.534, 0.12, 0.58],
    "metadata": {
        "product_name": "Studio Monitor Headphones",
        "category": "electronics",
        "brand": "Audio-Technica",
        "price": 149.99
    }
}
],
"limit": 2,
"offset": 0,
"count": 2
}

```

For more information about the PL/SQL implementation, including parameters, see `LIST_VECTORS`.

Delete Vectors

Use the `delete_vectors` operation to delete vectors from a table by their IDs.

The operation removes the specified vectors and their associated metadata from the table, essentially deleting the rows from the table for the given IDs.

An error is raised if the specified table does not exist.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of using `delete_vectors`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#delete specific vectors

```

```

response = client.delete_vectors(
    table_name='products',
    ids=['prod_old_1', 'prod_old_2', 'prod_old_3']
)

print(response)

```

A JSON response is returned confirming successful deletion with a count of deleted vectors.

Example response:

```

{
  "message": "Deleted 3 vector records"
}

```

Curl

See how POST `/vecdb/vector-tables/{vector_table_name}/delete` can be used in the following example:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/myvectors/delete" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "ids": ["vec1", "vec2"]
  }'

```

Responses:

- Example 200 response:

```

{
  "message": "Deleted 2 rows successfully"
}

```

- 400- the request body included invalid parameters.
- 404 - the specified vector table does not exist.

For more information, see [REST API Reference](#).

PL/SQL

See how `DBMS_VECTOR_DATABASE.DELETE_VECTORS` can be used in the following example:

```
dbms_vector_database.delete_vectors(  
    'myvectors',  
    JSON('["vec1", "vec2"]')  
);
```

Example output:

```
{  
  "message": "Deleted 2 rows successfully"  
}
```

For more information about the PL/SQL implementation, including parameters, see `DELETE_VECTORS`.

List Vector Load Jobs

Use the `list_vector_load_jobs` operation to retrieve a list of all bulk load operations.

The operation returns metadata for all load jobs including their states and progress.

- [Python](#)
- [Curl](#)

Python

See the following for examples of how to implement `list_vector_load_jobs`:

```
from oracle_vecdb import OracleVecDB, Configuration  
  
client = OracleVecDB(Configuration(  
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",  
    # or username="<user>", password="<pass>"  
    access_token="<bearer-token>",  
))  
  
load_jobs = client.list_vector_load_jobs()  
print(load_jobs)
```

A JSON response is returned containing an array of load jobs with job names, owners, states, and timestamps.

Example response:

```
{  
  "items": [  

```

```

{
  "job_name": "VECDB_LOAD_ABC123",
  "job_creator": "VECTOR3",
  "operation": "LOAD",
  "state": "SUCCEEDED",
  "links": [
    {
      "rel": "self",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/load/
jobs/vecdb_load_abc123"
    },
    {
      "rel": "related",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/load/
jobs/vecdb_load_abc123/jobfile"
    }
  ]
},
"count": 1
}

```

For more information about the `list_vector_load_jobs` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/load/jobs/` can be used in the following example:

```

curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/load/jobs/" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"

```

Responses:

- Example 200 response:

```

{
  "count": 1,
  "hasMore": false,
  "items": [
    {
      "job_creator": "APPUSER",
      "job_name": "VECDB_LOAD_20260501110000",
      "job_type": "PLSQL_BLOCK",
      "operation": "LOAD",
      "state": "SUCCEEDED",
      "start_date": "2026-05-01T11:00:00.118886Z",
      "links": [

```

```

        {
            "href": "/vecdb/load/jobs/vecdb_load_20260501110000",
            "rel": "self"
        },
        {
            "href": "/vecdb/load/jobs/vecdb_load_20260501110000/jobfile",
            "rel": "related"
        }
    ]
}
],
"limit": 25,
"offset": 0
}

```

For more information, see [REST API Reference](#).

Describe a Bulk Load Job

Use the `describe_vector_load_job` operation to get the status of a bulk load operation.

The operation returns details about an asynchronous load job initiated by a `load_vectors` operation.

-
- [Python](#)
 - [Curl](#)

Python

See the following for examples of how to implement `describe_vector_load_job`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#initiate a load job using load_vectors
response = client.load_vectors(
    table_name='products',
    url='https://objectstorage.region.oraclecloud.com/.../vectors.csv',
    params={'credential': 'OCI_CREDENTIAL'})

#monitor the load progress
status = client.describe_vector_load_job(response.job_name)
print(status)

```

A JSON response is returned containing job status, progress, and statistics.

Example response:

```
{
  "job_name": "VECDDB_LOAD_ABC123",
  "job_creator": "VECTOR3",
  "operation": "LOAD",
  "state": "SUCCEEDED",
  "links": [
    {
      "rel": "collection",
      "href": "https://<host>/ords/<schema>/_/db-api/stable/vecddb/load/jobs/"
    },
    {
      "rel": "self",
      "href": "https://<host>/ords/<schema>/_/db-api/stable/vecddb/load/jobs/vecddb_load_abc123/"
    },
    {
      "rel": "related",
      "href": "https://<host>/ords/<schema>/_/db-api/stable/vecddb/load/jobs/vecddb_load_abc123/jobfile"
    }
  ]
}
```

For more information about the `describe_vector_load_job` operation, see [Python API Reference](#).

Curl

See how `GET /vecddb/load/jobs/{load_job_name}/` can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_/db-api/stable/vecddb/load/jobs/VECDDB_LOAD_20260501110000/" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Example 200 response:

```
{
  "job_creator": "APPUSER",
  "job_name": "VECDDB_LOAD_20260501110000",
  "job_type": "PLSQL_BLOCK",
  "operation": "LOAD",
  "state": "SUCCEEDED",
  "start_date": "2026-05-01T11:00:00.11886Z",
  "links": [
    {
```

```

        "href": "/vecdb/load/jobs/",
        "rel": "collection"
    },
    {
        "href": "/vecdb/load/jobs/vecdb_load_20260501110000/",
        "rel": "self"
    },
    {
        "href": "/vecdb/load/jobs/vecdb_load_20260501110000/jobfile",
        "rel": "related"
    }
]
}

```

For more information, see [REST API Reference](#).

Retrieve Log Output for Bulk Loads

Use the `get_vector_load_job_log` operation to retrieve the log output for a bulk load job.

The operation returns log file metadata and contents.

-
- [Python](#)
 - [Curl](#)

Python

See the following for examples of how to implement `get_vector_load_job_log`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#initiate a load job using load_vectors
response = client.load_vectors(
    table_name='products',
    url='https://objectstorage.region.oraclecloud.com/.../vectors.csv',
    params={'credential': 'OCI_CREDENTIAL'})

#retrieve the log output
log_output = client.get_vector_load_job_log(response.job_name)
print(log_output)

```

A JSON response is returned containing log file metadata and contents.

For more information about the `get_vector_load_job_log` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/load/jobs/{load_job_name}/jobfile` can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/load/jobs/
  VECDB_LOAD_20260501110000/jobfile" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  -H "Authorization: Bearer <access_token>"
```

Example 200 response:

```
{
  "actual_start_date": "2026-05-01T11:00:00.22485Z",
  "additional_info": null,
  "error#": 0,
  "job_name": "VECDB_LOAD_20260501110000",
  "links": [
    {
      "href": "/vecdb/load/jobs/VECDB_LOAD_20260501110000/",
      "rel": "collection"
    }
  ],
  "log_date": "2026-05-01T11:00:00.572103Z",
  "run_duration": "P0DT0H0M0S",
  "status": "SUCCEEDED"
}
```

For more information, see [REST API Reference](#).

Search Data

Perform semantic search, including with advanced options and filtering, using the included operations. Make sure the output is useful and relevant to your use-case by reranking search results.

- [Semantic Search](#)
- [Filter by Metadata](#)
- [Search with Advanced Options](#)
- [Rerank Results](#)

Semantic Search

Use the query operation to search a vector table using a query vector, text, or record ID.

The operation performs similarity search to find the most similar vectors in the given table. Filtering by metadata and by various distance metrics is supported. The API retrieves IDs, metadata, vectors, and similarity scores of the most similar items from the table.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following example of semantic search using query:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#search by query vector
vector_results = client.query(
    table_name='products',
    query_by={'vector': [0.1, 0.2, 0.3, ...]},
    top_k=10
)
print(vector_results.items[0].id)

#search by text
text_results = client.query(
    table_name='products',
    query_by={'text': 'wireless headphones'},
    top_k=5
)
print(text_results)

#find similar items to an existing product
similar_item_results = client.query(
    table_name='products',
    query_by={'id': 'prod_12345'},
    top_k=10,
    filters={'category': {'$eq': 'electronics'}}
)
print(similar_item_results)
```

A JSON response is returned, containing the following information:

- An array of matching results with IDs, distances, and metadata
- Distance scores (a lower value is more similar for most metrics)
- Vector values (if `include_vectors` is set to `true`)

Example response:

```
{
  "items": [
    {
      "id": "prod_001",
      "metadata": {
        "name": "Aurora Trail Boots",
        "category": "footwear",
        "price": 129.99,
        "color": "midnight blue"
      },
      "vector": null,
      "distance": 0.0
    },
    {
      "id": "prod_008",
      "metadata": {
        "name": "Glacier Insulated Bottle",
        "category": "accessories",
        "price": 34.0,
        "color": "arctic white"
      },
      "vector": null,
      "distance": 0.00393
    }
  ]
}
```

For more information about the query operation, see [Python API Reference](#).

Curl

See how `POST /vecdb/vector-tables/{vector_table_name}/query` can be used in the following examples:

- Search by query vector:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/product_vectors/query" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  -H "Authorization: Bearer <access_token>" \
```

```
-d '{
  "queryBy": {"vector": [0.1, 0.5, 0.4, 0.56]},
  "topK": 5,
  "includeVectors": false
}'
```

- Search by text with filtering:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/
product_vectors/query" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  -u "<user>:<password>" \
  -d '{
    "queryBy": {"text": "good headphone"},
    "topK": 5,
    "filters": {
      "$and": [
        {"category": {"$eq": "electronics"}},
        {"price": {"$lt": 200}}
      ]
    },
    "outputSelector": ["product_name", "price", "product_description"],
    "includeVectors": false
  }'
```

- Find similar items to an existing product:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/
product_vectors/query" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  -u "<user>:<password>" \
  -d '{
    "queryBy": {"id": "id2"},
    "topK": 10,
    "filters": {"category": {"$eq": "electronics"}}
  }'
```

- Search with custom distance metric:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-tables/
product_vectors/query" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  -u "<user>:<password>" \
  -d '{
    "queryBy": {"text": "good headphone"},
    "topK": 10,
    "advancedOptions": {
      "distance_metric": "EUCLIDEAN",
      "idx_parameters": {
        "efsearch": 128,

```

```

        "neighbor_part": 4
      },
      "advanced_params": {
        "rescore_factor": 2
      }
    },
    "includeVectors": true
  }'

```

Responses:

- Example 200 response:

```

{
  "results": [
    {
      "id": "id2",
      "distance": 0.212,
      "metadata": {
        "product_name": "Studio Monitor Headphones",
        "price": 149.99,
        "product_description": "Professional wired headphones with
accurate sound for studio use."
      }
    },
    {
      "id": "id5",
      "distance": 0.678,
      "metadata": {
        "product_name": "Wireless Earbuds",
        "price": 199.99,
        "product_description": "Noise-cancelling wireless earbuds with
rich sound and secure fit."
      }
    }
  ]
}

```

- 400 - the request body included invalid parameters.
- 404 - the specified vector table does not exist.

For more information, see [REST API Reference](#).

PL/SQL

See how `DBMS_VECTOR_DATABASE.SEARCH` can be used in the following example:

```

dbms_vector_database.search(
  table_name      => 'product_vectors',
  query_by        => JSON('{"text": "good headphone"}'),
  top_k           => 5,
  filters         => JSON('{
    "$and": [
      { "category": { "$eq": "electronics" } },
      { "price":    { "$lt": 200 } }
    ]
  }')
)

```

```

    }'),
    include_vectors => false
);

```

Example response:

```

{"result" :
[
  {"id":"id2", "vector_distance" : 0.212,
   "metadata":{"product_name":"Studio Monitor Headphones","price":149.99,
    "product_description":"Professional wired headphones with
accurate sound for studio use."}},
  {"id":"id5", "vector_distance" : 0.678,
   "metadata":{"product_name":"Wireless Earbuds",,"price":199.99,
    "product_description":"Noise-cancelling wireless earbuds with rich sound
and secure fit."}}
]
}

```

For more information about the PL/SQL implementation, including parameters, see [SEARCH](#).

Filter by Metadata

Optionally use the `filters` parameter of the query API to provide metadata filters to narrow search results.

For details and the full list of parameters accepted by the query API, see [Semantic Search](#).

The `filters` parameter is used to filter search results. The following operators are supported:

Operator	Input Type	Description	Example
\$eq	String, Number, Boolean	Equal to operator for single values Retrieves records where a field value equals the specified value.	<pre>{ "genre": { "\$eq": "documentary" } }</pre>
\$ne	String, Number, Boolean	Not equal to operator Retrieves records where a field value does not equal the specified value.	<pre>{ "genre": { "\$ne": "drama" } }</pre>

Operator	Input Type	Description	Example
\$gt	Number	Greater than operator Retrieves records where a field value is greater than the specified value.	{ "year": { "\$gt": 2019 } }
\$gte	Number	Greater than or equal to operator Retrieves records where a field value is greater than or equal to the specified value.	{ "year": { "\$gte": 2020 } }
\$lt	Number	Less than operator Retrieves records where a field value is less than the specified value.	{ "year": { "\$lt": 2020 } }
\$lte	Number	Less than or equal to operator Retrieves records where a field value is less than or equal to the specified value.	{ "year": { "\$lte": 2020 } }
\$in	Non-empty array of primitives	In an array operator Retrieves records where a field value corresponds to any value in the specified array.	{ "genre": { "\$in": ["comedy", "documentary"] } }

Operator	Input Type	Description	Example
\$nin	Non-empty array of primitives	Not in an array operator Retrieves records where a field value does not correspond to any values in the specified array.	<pre>{ "genre": { "\$nin": ["comedy", "documentary"] } }</pre>
\$exists	Boolean	Field existence operator Retrieves records that contain (if set to True) or do not contain (if set to False) a specified value.	<pre>{ "genre": { "\$exist": true } } { "genre": { "\$exist": false } }</pre>
\$and	Non-empty array of filters	Logical AND operator Retrieves records that contain all specified values or conditions.	<pre>{ "\$and": [{ "genre": { "\$eq": "drama" } }, { "year": { "\$gte": 2020 } }] }</pre>

Operator	Input Type	Description	Example
\$or	Non-empty array of filters	Logical OR operator Retrieves records that contain at least one of the specified values or conditions.	<pre>{ "\$or": [{ "genre": { "\$eq": "drama" } }, { "year": { "\$gte": 2020 } }] }</pre>

See the following for an example of a semantic search using filters to specify that results must have a category equal to electronics and a price less than 200. This example uses the Python query API:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#search by text with filtering
text_filter_results = client.query(
    table_name='products',
    query_by={'text': 'wireless headphones'},
    top_k=5,
    filters={
        '$and': [
            {'category': {'$eq': 'electronics'}},
            {'price': {'$lt': 200}}
        ]
    }
)
print(text_filter_results)
```

Examples with filters using Curl and PL/SQL can be found at [Semantic Search](#).

Search with Advanced Options

Optionally use the `advanced_options` parameter of the query operation to further customize semantic searches.

For details about implementing the query operation using Python, Curl, and PL/SQL, see [Semantic Search](#).

The `advanced_options` (`advancedOptions` when using a Curl request) parameter can be used to provide search tuning parameters. The following are supported:

- `distance_metric`: Optionally used to override the default metric with a metric of your choice. The supported metrics are as follows:

- COSINE (default)
- MANHATTAN
- HAMMING
- JACCARD
- DOT
- EUCLIDEAN
- L2_SQUARED
- EUCLIDEAN_SQUARED

For more information about the available distance metrics, see Vector Distance Metrics in *Oracle Database AI Vector Search User's Guide*.

- `accuracy`: Trade-off between speed and accuracy, provided as a number from 0 to 100.

More information about target accuracy can be found at Understand Approximate Similarity Search Using Vector Indexes in *Oracle Database AI Vector Search User's Guide*.

- `idx_parameters`: Index-specific parameters. For example, the following values can be specified:

```
"idx_parameters" : {  
  "efsearch" : <number>,  
  "neighbor partition probes" : <number>  
}
```

The `efsearch` parameter is used to impose a certain maximum number of candidates to be considered while probing a vector index.

The `neighbor partition probes` parameter is used to impose a certain maximum number of partitions to be probed by the search.

See the following for a Python example of a semantic search using query with `advanced_options` to specify a distance metric and target accuracy:

```
from oracle_vecdb import OracleVecDB, Configuration  
  
client = OracleVecDB(Configuration(  
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",  
    access_token="<bearer-token>", # or username="<user>", password="<pass>"  
))
```

```
#search using a specified distance metric
distance_metric_results = client.query(
    table_name='products',
    query_by={'text': 'laptop'},
    top_k=10,
    advanced_options={
        'distance_metric': 'EUCLIDEAN',
        'accuracy': 95,
        'idx_parameters': {
            'efsearch': 128,
            'neighbor partition probes': 4
        }
    },
    include_vectors=True
)

print(distance_metric_results)
```

Example response:

```
{
  "items": [
    {
      "id": "prod_001",
      "metadata": {
        "name": "Aurora Trail Boots",
        "category": "footwear",
        "price": 129.99,
        "color": "midnight blue"
      },
      "vector": null,
      "distance": 0.0
    },
    {
      "id": "prod_008",
      "metadata": {
        "name": "Glacier Insulated Bottle",
        "category": "accessories",
        "price": 34.0,
        "color": "arctic white"
      },
      "vector": null,
      "distance": 0.00393
    }
  ]
}
```

Rerank Results

Use the rerank operation to re-rank search results based on relevance to a query.

The operation uses a reranking model to score and reorder documents relative to a query. This improves search quality by performing a more detailed comparison between the query and each candidate document.

For details about loading your own reranking model (in ONNX format) into the database, see [Load Your Own ONNX Model](#).

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of how to rerank query search results using `rerank()`.

Begin by loading a reranking model into your database. This action only needs to be completed once.

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_/db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#first, load a reranking model into your database
load_mod_response = client.load_model(
    model_name='cohere-rerank-3.5',
    url='https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace/b/
bucket/o/model.onnx',
    model_params={
        'provider': 'database',
        'credential': 'OCI_CREDENTIAL'
    }
)
print(load_mod_response)

#verify that the model was loaded by using the list_models() API
models = client.list_models()
print(models)
```

Now that the model is successfully loaded, you can perform a search and rerank the results.

```
#perform an initial search
import json
search_results = client.query(
    table_name='documents',
    query_by={'text': 'machine learning'},
    top_k=20
)
print(search_results)

#then, rerank the results for better relevance]
```

```

reranked = client.rerank(
    query='machine learning applications in healthcare',
    documents=[
        result.metadata["content"]
        for result in search_results.items
        if result.metadata and "content" in result.metadata
    ],
    model_name='cohere-rerank-3.5',
    model_params={'top_n': 5}
)
for item in reranked.items:
    original = search_results.items[item.index]
    print(f"{item.score:.3f} - {original.metadata['content']}")

```

A JSON response is returned containing reranked documents with relevance scores. An error is raised if the model is not loaded or if inputs are invalid.

Example response:

```

{
  "items": [
    {
      "index": 0,
      "score": 0.82
    }
  ]
}

```

For more information about the rerank operation, see [Python API Reference](#).

Curl

See how POST `/vecdb/rerank` can be used in the following example:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/rerank" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "modelName": "RERANK_MODEL",
    "query": "good headphone",
    "documents": [
      "Studio Monitor Headphones",
      "Wireless Earbuds",
      "Coffee Grinder"
    ],
    "modelParams": {

```

```

        "top_n": 2
    }
}'

```

Responses:

- Example 200 response:

```

[
  {
    "index": 0,
    "score": 0.92
  },
  {
    "index": 1,
    "score": 0.76
  }
]

```

- 400 - the request body included invalid parameters.

For more information about POST /vecdb/rerank, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.RERANK can be used in the following example:

```

dbms_vector_database.rerank(
    model_name => 'RERANK_MODEL',
    query => 'good headphone',
    documents => JSON('[
        "Studio Monitor Headphones",
        "Wireless Earbuds",
        "Coffee Grinder"
    ]'),
    model_params => JSON('{"top_n": 2}')
);

```

Example response:

```

[
  {
    "index": 0,
    "score": 0.92
  },
  {
    "index": 1,
    "score": 0.76
  }
]

```

For more information about the PL/SQL implementation, including parameters, see RERANK.

Manage Indexes

Use the included operations to manage vector indexes in your Vector Database.

- [Describe an Index](#)
- [Create an Index](#)
- [Rebuild an Index](#)
- [Drop an Index](#)
- [List Vector Load Jobs](#)
- [Describe an Index Job](#)
- [Retrieve an Index Job Log](#)

Describe an Index

Use the `describe_index` operation to get the current status and configuration of a vector table's index.

The operation returns detailed information about the index including its type, parameters, build status, and statistics.

An error is raised if the specified table or index does not exist.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of using `describe_index`:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

status = client.describe_index(table_name='products')

print(status)
```

A JSON response is returned containing the following information:

- Index type (IVF, HNSW)
- Built status (BUILDING, READY, FAILED)
- Index parameters

- Statistics (indexed vectors, memory usage)

For more information about the `describe_index` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/vector-indexes/{vector_table_name}` can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-indexes/
product_vectors" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Responses:

- Example 200 response:

```
{
  "Index Status": "Index build in progress...\nCurrent progress: 75%"
}
```

- 400 - the request body included invalid parameters.
- 404 - the vector table was not found.

For more information about `GET /vecdb/vector-indexes/{vector_table_name}`, see [REST API Reference](#).

PL/SQL

See how `DBMS_VECTOR_DATABASE.INDEX_BUILD_STATUS` can be used in the following example:

```
dbms_vector_database.index_build_status('product_vectors');
```

Example response:

```
{
  "Index Status": "Index build in progress...\nCurrent progress: 75%"
}
```

For more information about the PL/SQL implementation, including parameters, see [INDEX_BUILD_STATUS](#).

Create an Index

Use the `create_index` operation to create a vector index on a table to enable fast similarity search.

The operation creates an index for efficient approximate nearest neighbor (ANN) search. The index creation runs asynchronously as a background job.

Both Inverted File Flat (IVF) and Hierarchical Navigable Small World (HNSW) indexes are supported.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example of using `create_index`.

```
describe_indexdescribe_index_job

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

#create an index with default IVF settings
def_ivf_response = client.create_index(table_name='products')
print(def_ivf_response)

#create an HNSW index with custom parameters
hnsw_response = client.create_index(
    table_name='products',
    index_params={
        'vector_index_params': {
            'auto_index': True,
            'organization': 'INMEMORY GRAPH',
            'distance_metric': 'COSINE',
            'quantization_type': 'SCALAR',
            'compression_ratio': 4,
            'distribute_params': {
                'distribute_method': 'AUTO'
            },
        },
        'advanced_params': {
            'neighbors': 32,
            'efConstruction': 200,
            'rescore_factor': 10,
            'algorithm': 'uniform_quantization'
        }
    }
)
```

```

        },
        'metadata_index_params': {
            'auto_index': True,
            'include_paths': ['tenant', 'category'],
            'exclude_paths': ['body']
        },
        'parallel_creation': 4
    }
)
print(hnsw_response)

#Create an IVF index with explicit defaults
ex_ivf_response = client.create_index(
    table_name='products',
    index_params={
        'vector_index_params': {
            'organization': 'PARTITIONS',
            'distance_metric': 'COSINE',
            'advanced_params': {
                'partitions': 16
            }
        }
    }
)
print(ex_ivf_response)

#monitor index creation progress
status = client.describe_index(table_name='products')
print(status)

```

A JSON response is returned containing the index job ID and status.

For more information about the `create_index` operation, see [Python API Reference](#).

Curl

See how `POST /vecdb/vector-indexes/` can be used in the following examples.

- Create an HNSW index with metadata index paths:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "tableName": "product_vectors",

```

```

    "indexParams": {
      "vector_index_params": {
        "auto_index": true,
        "organization": "INMEMORY GRAPH",
        "distance_metric": "COSINE",
        "accuracy": 95,
        "quantization_type": "SCALAR",
        "compression_ratio": 4,
        "advanced_params": {
          "neighbors": 32,
          "efConstruction": 128,
          "rescore_factor": 4,
          "algorithm": "uniform_quantization"
        }
      },
      "metadata_index_params": {
        "auto_index": true,
        "include_paths": ["category", "price"]
      },
      "parallel_creation": 2
    }
  }
}'

```

- Create an IVF index with explicit defaults:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  -u "<user>:<password>" \
  -d '{
    "tableName": "product_vectors",
    "indexParams": {
      "vector_index_params": {
        "auto_index": true,
        "organization": "PARTITIONS",
        "distance_metric": "COSINE",
        "advanced_params": {
          "partitions": 16
        }
      },
      "parallel_creation": 2
    }
  }'

```

Responses:

- Example 200 response:

```

{
  "job_creator": "APPUSER",
  "job_name": "VECDDB_CREATE_INDEX_20260501100000",
  "job_type": "PLSQL_BLOCK",
  "operation": "CREATE",
  "state": "SUCCEEDED",

```

```

"start_date": "2026-05-01T10:00:00.302367Z",
"links": [
  {
    "href": "/vecdb/vector-indexes/jobs/",
    "rel": "collection"
  },
  {
    "href": "/vecdb/vector-indexes/jobs/vecdb_create_index_20260501100000/",
    "rel": "self"
  },
  {
    "href": "/vecdb/vector-indexes/jobs/vecdb_create_index_20260501100000/jobfile",
    "rel": "related"
  }
]
}

```

- 400 - the request body included invalid parameters.
- 404 - the vector table was not found.

For more information about POST /vecdb/vector-indexes/, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.CREATE_INDEX can be used in the following example:

```

dbms_vector_database.create_index(
  table_name => 'product_vectors',
  index_params => JSON('{
    "vector_index_params": {
      "auto_index": true,
      "organization": "INMEMORY GRAPH",
      "distance_metric": "COSINE",
      "accuracy": 95,
      "quantization_type": "SCALAR",
      "compression_ratio": 4,
      "advanced_params": {
        "neighbors": 32,
        "efConstruction": 128,
        "rescore_factor": 4,
        "algorithm": "uniform_quantization"
      }
    },
    "metadata_index_params": {
      "auto_index": true,
      "include_paths": ["category", "price"]
    },
    "parallel_creation": 2
  }')
);

```

Example response:

```
{
  "message": "Vector index VECIDX_PRODUCT_VECTORS_20260501T100000 and
Metadata index(es) for paths [category, price] created successfully for table
PRODUCT_VECTORS"
}
```

For more about the PL/SQL implementation, including parameters, see `CREATE_INDEX`.

Rebuild an Index

Use the `rebuild_index` operation to rebuild an existing vector index with updated parameters.

The operation recreates the existing index, optionally with new configuration parameters. This is useful for optimizing search performance or adjusting to changed data distributions. The rebuild runs asynchronously as a background job.

- [Python](#)
- [Curl](#)
- [PL/SQL](#)

Python

See the following for an example of an index rebuild. In this case, all indexes are rebuilt with parallel DDL:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

response = client.rebuild_index(
    table_name='products',
    index_params={
        'index_type': 'all',
        'parallel_creation': 4
    }
)

print(response)
```

A JSON response is returned containing the rebuild job ID and status.

For more information about the `rebuild_index` operation, see [Python API Reference](#).

Curl

See how POST /vecdb/vector-indexes/{vector_table_name} can be used in the following examples:

- Vector-only rebuild:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/product_vectors" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "indexParams": {
      "index_type": "vector",
      "parallel_creation": 4
    }
  }'
```

- Metadata-only rebuild:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/product_vectors" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "indexParams": {
      "index_type": "metadata",
      "metadata_index_params": {
        "include_paths": ["tenant"]
      }
    }
  }'
```

- All-index rebuild:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/product_vectors" \
```

```

-H "Content-Type: application/json" \
-H "Accept: application/json" \
# Choose ONE authentication method:

# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
# -H "Authorization: Bearer <access_token>" \

-d '{
  "indexParams": {
    "index_type": "all",
    "parallel_creation": 4
  }
}'

```

Responses:

- Example 200 response:

```

{
  "job_creator": "APPUSER",
  "job_name": "VECDB_REBUILD_INDEX_20260501103000",
  "job_type": "PLSQL_BLOCK",
  "operation": "REBUILD",
  "state": "RUNNING",
  "start_date": "2026-05-01T10:30:00.118886Z",
  "links": [
    {
      "href": "/vecdb/vector-indexes/jobs/",
      "rel": "collection"
    },
    {
      "href": "/vecdb/vector-indexes/jobs/vecdb_rebuild_index_20260501103000/",
      "rel": "self"
    },
    {
      "href": "/vecdb/vector-indexes/jobs/vecdb_rebuild_index_20260501103000/jobfile",
      "rel": "related"
    }
  ]
}

```

- 400 - the request body included invalid parameters.
- 404 - the vector table was not found.

For more information about POST `/vecdb/vector-indexes/{vector_table_name}`, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.REBUILD_INDEX can be used in the following example:

```
dbms_vector_database.rebuild_index(  
    table_name => 'product_vectors',  
    index_params => JSON('{  
        "vector_index_params": {  
            "auto_index": true,  
            "organization": "INMEMORY GRAPH",  
            "distance_metric": "COSINE",  
            "advanced_params": {  
                "neighbors": 32,  
                "efConstruction": 128,  
                "rescore_factor": 4  
            }  
        },  
        "parallel_creation": 2  
    }')  
);
```

Example response:

```
{  
  "message": "Vector index VECIDX_PRODUCT_VECTORS_20260501T100000 Rebuild  
successfully for table PRODUCT_VECTORS"  
}
```

For more information about the PL/SQL implementation, including parameters, see REBUILD_INDEX.

Drop an Index

Use the drop_index operation to drop the vector index from a table.

The operation removes the index while preserving the table and its data. Queries will fall back to exact search until a new index is created. A table name is accepted as input.

-
- [Python](#)
 - [Curl](#)
 - [PL/SQL](#)

Python

See the following for an example using drop_index to drop an index and inspect the response.

```
from oracle_vecdb import OracleVecDB, Configuration  
  
client = OracleVecDB(Configuration(  

```

```

rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
access_token="<bearer-token>", # or username="<user>", password="<pass>"
))

result = client.drop_index(
    table_name="products",
    index_params={"index_type": "all"},
)

print(result)

```

A JSON response is returned confirming index deletion. An error is raised if the table or index does not exist.

Example response:

```

{
  "message": "Index drop request submitted."
}

```

For more information about the drop_index operation, see Python API Reference.

Curl

See how POST /vecdb/vector-indexes/{vector_table_name}/delete can be used in the following examples:

- Drop vector indexes only:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/product_vectors/delete" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>" \

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>" \

  -d '{
    "indexParams": {
      "index_type": "vector"
    }
  }'

```

- Drop selected metadata indexes:

```

curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/product_vectors/delete" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \

```

```
# Choose ONE authentication method:

# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
-H "Authorization: Bearer <access_token>" \

-d '{
  "indexParams": {
    "index_type": "metadata",
    "metadata_index_params": {
      "include_paths": ["tenant"]
    }
  }
}'
```

- Drop all indexes:

```
curl -X POST \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/product_vectors/delete" \
  -H "Content-Type: application/json" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

# Option 1: Basic authentication
-u "<user>:<password>" \

# Option 2: OAuth Bearer token
-H "Authorization: Bearer <access_token>" \

-d '{
  "indexParams": {
    "index_type": "all"
  }
}'
```

Responses:

- Example 200 response:

```
{
  "message": "MVI index(es) for paths [tenant] dropped successfully from
table PRODUCT_VECTORS"
}
```

- 400 - the request body included invalid parameters.
- 404 - the vector table was not found.

For more information about POST /vecdb/vector-indexes/{vector_table_name}/delete, see REST API Reference.

PL/SQL

See how DBMS_VECTOR_DATABASE.DROP_INDEX can be used in the following example:

```
dbms_vector_database.drop_index(  
    table_name => 'product_vectors',  
    index_params => JSON('{  
        "index_type": "metadata",  
        "metadata_index_params": {  
            "include_paths": ["category", "price"]  
        }  
    }')  
);
```

Example response:

```
{  
  "message": "MVI index(es) for paths [category, price] dropped successfully  
from table PRODUCT_VECTORS"  
}
```

For more information about the PL/SQL implementation, including parameters, see DROP_INDEX.

List Vector Load Jobs

Use the `list_index_jobs` operation to list asynchronous index jobs.

The operation returns metadata for all CREATE and REBUILD index operations, including their current states and progress.

- [Python](#)
- [Curl](#)

Python

See the following for examples of how to implement `list_index_jobs`:

```
from oracle_vecdb import OracleVecDB, Configuration  
  
client = OracleVecDB(Configuration(  
    rest_url="https://<host>/ords/<schema>/_/db-api/stable/vecdb/",  
    # or username="<user>", password="<pass>"  
    access_token="<bearer-token>",  
))  
  
index_jobs = client.list_index_jobs()  
print(index_jobs)
```

A JSON response is returned containing an array of index jobs.

Example response:

```
{
  "items": [
    {
      "job_name": "VECDB_CREATE_INDEX_ABC123",
      "job_creator": "VECTOR3",
      "operation": "CREATE",
      "state": "SUCCEEDED",
      "links": [
        {
          "rel": "self",
          "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/jobs/vecdb_create_index_abc123/"
        },
        {
          "rel": "related",
          "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/jobs/vecdb_create_index_abc123/jobfile"
        }
      ]
    }
  ],
  "count": 1
}
```

For more information about the `list_index_jobs` operation, see Python API Reference.

Curl

See how GET `/vecdb/vector-indexes/jobs/` can be used in the following example:

```
curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-indexes/
jobs/" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"
```

Responses:

- Example 200 response:

```
{
  "count": 1,
  "hasMore": false,
  "items": [
    {
```

```

        "job_creator": "APPUSER",
        "job_name": "VECDDB_CREATE_INDEX_20260501100000",
        "job_type": "PLSQL_BLOCK",
        "operation": "CREATE",
        "state": "SUCCEEDED",
        "start_date": "2026-05-01T10:00:00.302367Z",
        "links": [
            {
                "href": "/vecdb/vector-indexes/jobs/vecdb_create_index_20260501100000",
                "rel": "self"
            },
            {
                "href": "/vecdb/vector-indexes/jobs/vecdb_create_index_20260501100000/jobfile",
                "rel": "related"
            }
        ]
    },
    "limit": 25,
    "offset": 0
}

```

For more information about GET /vecdb/vector-indexes/jobs/, see REST API Reference.

Describe an Index Job

Use the `describe_index_job` operation to describe the current state of an index job.

The operation returns metadata and status information for a specific index build job. Details are returned about an asynchronous index creation or rebuild job, including its current state, progress, owner, and log file location.

-
- [Python](#)
 - [Curl](#)

Python

See the following for examples of how to implement `describe_index_job`:

```

from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    # or username="<user>", password="<pass>"
    access_token="<bearer-token>",
))

#create an index-build job

```

```

response = client.create_index(table_name='products')

#describe the index job
job_status = client.describe_index_job(response.job_name)
print(job_status)

```

A JSON response is returned describing the index job.

Example response:

```

{
  "job_name": "VECDB_CREATE_INDEX_ABC123",
  "job_creator": "VECTOR3",
  "operation": "CREATE",
  "state": "SUCCEEDED",
  "links": [
    {
      "rel": "collection",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/jobs/"
    },
    {
      "rel": "self",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/jobs/vecdb_create_index_abc123/"
    },
    {
      "rel": "related",
      "href": "https://<host>/ords/<schema>/_db-api/stable/vecdb/vector-
indexes/jobs/vecdb_create_index_abc123/jobfile"
    }
  ]
}

```

For more information about the `describe_index_job` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/vector-indexes/jobs/{index_job_name}/` can be used in the following example:

```

curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-indexes/
jobs/VECDB_CREATE_INDEX_20260501100000/" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"

```

Responses:

- Example 200 response:

```
{
  "job_creator": "APPUSER",
  "job_name": "VECDB_CREATE_INDEX_20260501100000",
  "job_type": "PLSQL_BLOCK",
  "operation": "CREATE",
  "state": "SUCCEEDED",
  "start_date": "2026-05-01T10:00:00.302367Z",
  "links": [
    {
      "href": "/vecdb/vector-indexes/jobs/",
      "rel": "collection"
    },
    {
      "href": "/vecdb/vector-indexes/jobs/vecdb_create_index_20260501100000/",
      "rel": "self"
    },
    {
      "href": "/vecdb/vector-indexes/jobs/vecdb_create_index_20260501100000/jobfile",
      "rel": "related"
    }
  ]
}
```

For more information about GET `/vecdb/vector-indexes/jobs/{index_job_name}/`, see REST API Reference.

Retrieve an Index Job Log

Use the `get_index_job_log` operation to retrieve the log output for an index build job.

The operation returns the detailed log file contents for diagnosing index creation issues or monitoring progress.

-
- [Python](#)
 - [Curl](#)

Python

See the following for examples of how to implement `get_index_job_log`:

```
from oracle_vecdb import OracleVecDB, Configuration

client = OracleVecDB(Configuration(
    rest_url="https://<host>/ords/<schema>/_db-api/stable/vecdb/",
    # or username="<user>", password="<pass>"
    access_token="<bearer-token>",
```

```

))

#create an index to initiate a load job
response = client.create_index(table_name='products')

#retrieve the log output
log_output = client.get_index_job_log(response.job_name)
print(log_output)

```

A JSON response is returned containing log file metadata and contents.

For more information about the `get_index_job_log` operation, see [Python API Reference](#).

Curl

See how `GET /vecdb/vector-indexes/jobs/{index_job_name}/jobfile` can be used in the following example:

```

curl -X GET \
  "https://<host>:<port>/ords/<schema>/_db-api/stable/vecdb/vector-indexes/
jobs/VECDDB_CREATE_INDEX_20260501100000/jobfile" \
  -H "Accept: application/json" \
  # Choose ONE authentication method:

  # Option 1: Basic authentication
  -u "<user>:<password>"

  # Option 2: OAuth Bearer token
  # -H "Authorization: Bearer <access_token>"

```

Responses:

- Example 200 response:

```

{
  "actual_start_date": "2026-05-01T10:00:00.32485Z",
  "additional_info": null,
  "error#": 0,
  "job_name": "VECDDB_CREATE_INDEX_20260501100000",
  "links": [
    {
      "href": "/vecdb/vector-indexes/jobs/
VECDDB_CREATE_INDEX_20260501100000/",
      "rel": "collection"
    }
  ],
  "log_date": "2026-05-01T10:00:00.772103Z",
  "run_duration": "P0DT0H0M0S",
  "status": "SUCCEEDED"
}

```

For more information about `GET /vecdb/vector-indexes/jobs/{index_job_name}/jobfile`, see [REST API Reference](#).

Manage Vector Database Users

Topics:

- [Create a New Vector Database User](#)
- [Grant Privileges to an Existing Vector Database User](#)

Create a New Vector Database User

Complete the following steps to create a new Vector Database user.

1. On the **Database Users** page within Database Actions, click **Create User**.
2. Complete the required fields for the new Vector Database user.
 - a. Assign the following roles and privileges based on your deployment type:
 - **For on-premises deployments:**
 - Select **Default** for the DB_DEVELOPER_ROLE role.
 - The CREATE MINING MODEL privilege must be granted after you create the user.
 - **For ADB-S:**
 - Select **Default** for the DB_DEVELOPER_ROLE role.
 - Assign the OML_DEVELOPER role.
 - Assign the OCI\$RESOURCE_PRINCIPAL privilege.
 - b. Select **Enable REST** to enable REST for the user.
 - c. For ADB-S, select **OML** to enable Oracle Machine Learning for the user.

For more information about the required fields, see [The Database Users Page](#).

3. Click **Create**.

This creates a new Vector Database user and grants the roles and privileges selected during user creation.

4. For on-premises deployments, grant the CREATE MINING MODEL privilege to the newly created user by executing the following SQL in Database Actions:

```
GRANT CREATE MINING MODEL TO <VECTOR_USER>;
```

5. For ADB-S, enable resource principal for the newly created user by executing the following SQL within Database Actions:

```
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(username =>  
'<VECTOR_USER>');
```

Note

Step 4 is necessary for the new user to access the vector database console.

Create a New Vector Database User Using SQL

Complete the following steps to create a new Vector Database user using SQL.

1. In the **Database actions** dropdown list, select **SQL** to use a SQL Worksheet.
2. Run the following SQL statements:

```
-- Create the user
create user <VECTOR_USER> identified by "<PASSWORD>";

-- For on-premises deployments, add these roles
grant db_developer_role to <VECTOR_USER>;
grant create mining model to <VECTOR_USER>;

-- For ADB-S, add these roles
grant db_developer_role, oml_developer, oci$resource_principal to
<VECTOR_USER>;

-- Enable REST
exec ords_admin.enable_schema(p_schema => '<VECTOR_USER>');

-- Set Quota
ALTER USER <VECTOR_USER> QUOTA UNLIMITED ON data;

--For ADB-S: Enable resource principal
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(username =>
'<VECTOR_USER>');
```

Grant Privileges to an Existing Vector Database User

Granting privileges to an existing user involves:

- Granting user roles
- Enabling the schema for ORDS
- Setting a table space quota
- Enabling resource principal for the user

Use the following SQL to prepare the user and enable resource principal:

```
-- Add Roles (safe if already granted)

--For On-premises deployments
grant db_developer_role, create mining model to <VECTOR_USER>;
--For ADB-S
grant db_developer_role, oml_developer, oci$resource_principal to
<VECTOR_USER>;

-- Enable REST (safe to re-run)
begin
  ords_admin.enable_schema(p_schema => '<VECTOR_USER>');
end;
/
```

```
-- Set Quota (overwrites existing quota)
ALTER USER <VECTOR_USER> QUOTA UNLIMITED ON data;

--For ADB-S: Enable Resource Principal (safe to re-run)
BEGIN
  DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(username => '<VECTOR_USER>');
END;
/
```

Configure an OAuth Client

Using Database Actions with OAuth 2.0-based authentication, you can ensure that your Autonomous AI Vector Database is accessed only by specific users or clients.

OAuth 2.0 is a standard Internet protocol that defines flows to provide conditional and limited access to a RESTful API. For more information, see *OAuth-Based Authentication in Oracle REST Data Services Developer's Guide*.

To navigate to the OAuth Clients page, first navigate to the **Development** tab in Database Actions and then select **REST**. To create an OAuth client, see the following steps:

1. In the OAuth Clients page, select **Create OAuth Client**.
2. Enter the following fields. The fields with an asterisk (*) are mandatory:

Client Definition tab

- **Owner:** This is a read only field that should be your Vector Database user.
- **Grant type:** Select the authorization grant type. In this case, for a client associated with a vector database, select **CLIENT_CRED** from the drop down menu. For information about the available grant types, see *OAuth Flows in Oracle REST Data Services Developer's Guide*.
- **Name:** Name of the client.
- **Description:** Description of the purpose of the client.
- **Support Email:** Enter the email where end users can contact the client for support.
- (Optional) **Support URI:** Enter the URI where end users can contact the client for support. Example: `http://www.myclientdomain.com/support/`

Show code: Select this option to view the PL/SQL code equivalent of the Create OAuth Client panel. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create OAuth Client panel.

3. Select roles that should be associated with the client.

Navigate to the **Roles** tab, select `oracle.dbtools.auth.roles.builtin.VecDB` from the available roles, and click the single right arrow to assign it to the new client.

4. Click **Create**.

The OAuth Client registered is displayed on the OAuth Clients page.

A preview of the client secret is displayed. Copy the secret value as this is the only instance when it is shown. If you forget the client secret, you can use the Rotate Secret action to change it. See *Managing Secrets in Using Oracle Database Actions*.

The Client ID value represents the secret credential for the OAuth client. Click **Show/Hide**  to see the values.

Test the secured REST service endpoint using a REST client or the cURL command line tool.

For more general, non-Vector Database specific information about managing OAuth clients using Database Actions, see Managing OAuth Clients in *Using Oracle Database Actions*.

5. Use the client ID and client secret to request an OAuth token.

Use the following curl command, with values replaced with your own input:

```
curl -i -k --user <client_id>:<client_secret> --data  
"grant_type=client_credentials" https://example.com/ords/vecdb_user/oauth/  
token
```

The access token returned can be used to make requests to an Autonomous AI Vector Database endpoint. For example, you can make a request to the summary endpoint to find the number of tables, vectors, and models you have in your Vector Database, as in the following:

```
curl --location 'https://example.com/ords/vecdb_user/_/db-api/stable/vecdb/  
summary' \  
--header 'Authorization: Bearer <access_token>'
```

Part III

Administration

This part provides information about the following topics:

Topics:

- [The Database Users Page](#)
- [The APEX Workspaces Page](#)
- [The Data Pump Page](#)

10

The Database Users Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Database Users page enables you to perform user management tasks such as create, edit, enable REST, and delete, and create and manage self service schemas. For user management, the actions available are based on the user privileges (CREATE, ALTER, DROP) granted to you.


Users with no assigned privileges can still access the Database Users page to browse all users. However, the only action available to them is changing their password.

To navigate to the Database Users page, do either of the following:

- In the Launchpad page, select the **Administration** tab and click **Database Users**.
- Click **Selector**  to display the navigation menu. Under Administration, select **Database users**.

The Database Users page consists of two parts: Current User and All Users.

Current User



Displays information about the current user such as user name, whether REST Enabled or not, REST Alias, account expiration (in days), and the last login date and time. The icon on the left displays the user status with one of the following colours: green (Open), blue (Locked), and red (Expired).


The URL at the bottom is displayed only if the user is REST Enabled. It provides the URL to the Database Actions user login page. Use **Copy to Clipboard**  to copy the URL to the user's clipboard.

Click **Actions**  to open the context menu. The actions available are:

- **Edit**: Opens the Edit User Dialog, where you can edit current user information. See [Creating or Editing a User](#).

- **Enable REST:** Enables REST for a user where disabled. When this option is selected the first time, it opens the Enable REST dialog. See [Enabling REST](#).
- **Disable REST:** Disables REST where enabled for a user.
- **Drop REST Services:** Removes the REST data for a user, such as REST Alias, Base Path and so on, which is stored in Oracle REST Data Services (ORDS).
- **Delete:** Opens the Delete User dialog, where you can delete the user. See [Deleting a User](#).

All Users



Displays information about all other users that have been created in the database. You can use the Search field, which is case insensitive, to search for users, or sort the users in ascending or descending order using the sort icons, or filter by user status or REST status.

To create a user, click **Create User** to open the Create User dialog. For more information, see [Creating or Editing a User](#).

There are two views available:



(Card View): Displays the user information in a card view. This is the default display view. Each user card provides details such as user status, password expiry, user name, and the context menu.



(Grid View): Displays the user information in a tabular format. The last column in each row contains the context menu icon.

10.1 Creating or Editing a User

You can create a new database user or edit an existing database user.

- To create a new database user, click **Create User** in the Database Users page.
- To edit an existing database user, select **Edit** from the context menu for the associated user.

The user properties are grouped under two tabs: User and Granted Roles.

User Tab

Specifies general properties for the database user.

- **User Name:** The user name string. For an existing user, this field is read-only. To change the name, you must drop the user and create a new user with the desired name.
- **Password:** Password string for the new user, or new password for an existing user. You must also type the same password string for **Confirm Password**.

- **Password Expired:** If this option is selected, the password is marked as expired, and the user must change the password before being permitted to connect to the database.
- **Account is Locked:** If this option is selected, the user will not be permitted to connect to the database until a DBA user unlocks the account associated with this user.
- **Quota on tablespace USERS:** Enter or select to assign quota on the user's default tablespace. By default, the user has no quota on the default tablespace.
- **Web Access:** If this option is selected, the user is enabled for REST access. Expand **Web access advanced features** to specify the related fields: **Authorization required**, **REST Alias**, and **URL Mapping Type**.

Granted Roles Tab

Specifies roles to be granted to the user.

Use **Filter by role** to quickly locate the required roles.

For each role, you can select **Granted** to grant the role, **Admin** to permit the user to grant the role to other users, and **Default** to use the default settings for Granted and Admin. A new user is granted **CONNECT** and **RESOURCE** roles when Web Access is selected.

Show code: Select this option to view the PL/SQL code equivalent of the Create/Edit User slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create/Edit** in the Create/Edit User slider.

10.2 Enabling REST

You can enable REST for a user that has not been REST enabled.

In the Database Users page, select **Enable REST** from the user's context menu.

When you select Enable REST for a user for the first time, REST Enable User dialog is displayed. Subsequently, if you disable REST and then select Enable REST again, you receive a message stating that REST is enabled. In this case, the REST data previously provided is used for enabling REST. To enter new REST data, select **Drop REST Services** and then select **Enable REST** again.

Schema Alias: Enter the alias for the schema name to use in the URL.

The alias is entered at the time of signing in. See [Signing In to Database Actions](#)

URL Mapping Type: Select **BASE_PATH** or **BASE_URL**.

Authorization Required: For a schema, controls whether Oracle REST Data Services should require user authorization before allowing access to the Oracle REST Data Services metadata catalog of this schema.

10.3 Deleting a User

Use this option to delete users.

In the Database Users page, select **Delete** from the user's context menu to delete a user.

Note

The number of user's active sessions is displayed in the Delete User dialog window.

- **Cascade:** If this option is selected, all dependent objects are also deleted.
- **Drop REST Services:** If this option is selected, all user REST data is removed from ORDS.

Note

If you do not select this option when deleting a user, the next time you create a user with the same user name, it will still retain the REST-enabled property.

Click **Delete User** and a confirmation notification is displayed.

Note

An error notification is displayed, if the user has active sessions. In such cases, you must close the active sessions before you can delete the user.

10.4 Creating a Self Service Schema

Note

Available only if you have the **DBA** and **PDB_DBA** roles assigned.

The Self Service Schema page enables you to create a database schema on user request by generating a secure URL that is automatically sent by email to the user. Users can access a public page to request for a schema.

The steps are broken down into the following sections:

- Activate the Self Service Schema in Oracle REST Data Services (ORDS)
- Enable the Self Service Schema feature in Database Actions
- Request database schema by user
- Manage schema requests on the Self Service Schema Page

Activate the Self Service Schema in ORDS

1. By default, this feature is not available unless activated. To activate, set the configuration property in ORDS by using the command line or by adding the property to the XML configuration file.

- From the command line, use the following command:

```
ords config set feature.sdw.selfServiceSchema true
```

- In the configuration XML file, add the property as shown below:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE properties SYSTEM "http://java.sun.com/dtd/properties.dtd">
<properties>
<comment>Saved on Wed Nov 23 16:57:17 UTC 2022</comment>
```

```

<entry key="database.api.enabled">true</entry>
<entry key="feature.sdw.selfServiceSchema">true</entry>
<entry key="sdw.dev.doc.root">C:\Users\jsmith\Desktop\static-resources</entry>
<entry key="standalone.context.path">/ords</entry>
<entry key="standalone.http.port">8080</entry>
</properties>

```


After the ORDS configuration property is set to true, you can enable or disable the feature from the Database Users page in the Database Actions user interface.

Enable the Self Service Schema Feature in Database Actions

2. In the Database Users screen, click  **Self Service Schema Settings**, which is located next to Create User.
3. In the Self Service Schema settings slider:
 - a. Select **Enable Feature** to enable the Self Service Schema feature.
The URL to the Request Schema page appears, which can be shared with the user.
 - b. Click **Save**.

The **Self Service Schema** tab appears in the header next to User Management on the Database Users screen. Navigate to the Self Service Schema page using this tab.

Figure 10-1 Self Service Schema Tab




Request Database Schema by User

4. In the Request a Schema page, a user has to enter details such as Requested User Name, Full Name, Email, and Justification, click **Next** and then click **Submit**.

Note

Click  towards the top right corner in the Self Schema Service page (see [Figure 10-1](#)) to access the Request page URL.



The image shows a 'Request a Schema' form with a red database icon at the top. The form contains the following fields:

- Requested Username ***: Text input with 'DOEJ'.
- Full Name ***: Text input with 'Jane Doe'.
- Email ***: Text input with 'jane.doe@example.com'.
- Justification ***: Text area with 'Training'.


At the bottom left is a help icon (?), and at the bottom right is a blue 'Next' button.

A confirmation message appears. The user request becomes available on the Self Service Schema page in a card format.

Manage Schema Requests on the Self Service Schema Page

5. You can see and manage user requests on the Schema Requests section in the Self Service Schema page.

You can set up an SMTP server to automatically send emails to users to notify whether the schema request has been approved or rejected. To define SMTP settings, click  on the top right corner of the Schema Requests page. Enter fields such as Server Host, Server Port, Sender Name and Server Encryption to set up the server.




The image shows the 'Self Service Schema' page with the 'Schema Requests' section. It includes a search bar, filter, and sort options. A card for 'SELF_SERVICE_SCHEMA_II583' by 'John Smith (jsmith@abc.com)' is shown with a 'Pending' status and 'Approve'/'Reject' buttons. On the right, the 'Email settings' panel is visible with the following fields:

- Server host ***: Text input.
- Server port ***: Text input with '25' and up/down arrows.
- Sender ***: Text input with 'noreply@example.com'.
- Server encryption**: Dropdown menu with 'None' selected.

At the bottom right are 'Save' and 'Cancel' buttons.

6. After SMTP settings are set, you can approve or reject the request by clicking the appropriate button on the card.
 - If you click **Approve**, a panel with two tabs appears: Request Details and Initial Grants. You can make changes to the schema name and customize the initial roles for the requested schema. By default, the Resource and Connect roles are pre-selected. Click **Approve**.




- If you click **Reject**, you are prompted to specify a reason for rejection, which will be sent to the user.

The schema card is updated and shows the status of the request when it is approved or rejected.

Note


- If you cannot see the request card on the Schema page, check if any filters are set and clear them.
- After the schema is approved, you can still reject the schema. To do that, click the **Actions** icon on the card and select **Open Request Details**. Click **Reject**, enter a reason for rejection, and click **Submit**. The status of the card changes to Rejected and an email notification is sent to the user.

When you click **Approve**, a notification mail is sent to the user's email address.



7. Click **Get Started**. You are taken to a page where you need to set the password for the new schema. Enter **Password** and **Confirm Password** and click **Submit**.

8. Another mail is sent to the user notifying that the schema has been created and providing a **Start Using Schema** link. The user can use the link to log in to the schema after providing the new password.



After the schema is created, the status of the user request card on the Schema Requests page changes to **Fulfilled**.

Note

If you want to remove a schema request after the status changes to Fulfilled, you will need to use queries in the SQL worksheet.

11

The APEX Workspaces Page

Note

You can create and delete workspaces only if you have the **APEX_Administrator_Role** assigned.

You can create APEX workspaces using the APEX Workspaces page. For more information about APEX workspaces, see [Workspace and Application Administration](#) in the *Oracle Application Express Administration Guide*.

To navigate to this page, do either of the following:

- In the Launchpad page, select the **Administration** tab and click **APEX Workspaces**.
- Click **Selector**  to display the navigation menu. Under **Administration**, select **APEX Workspaces**.

For a specific workspace, you can perform the following actions using the  **Actions** context menu:

- **Create a Workspace:** See [Create a Workspace](#)
- **Delete a Workspace:** Removes the workspace and all associated applications from the database.
- **View Schema Details:** Navigates to the Database Users page to view the related schema information.

11.1 Create a Workspace

Note

You require Administrator privileges to create a workspace.

To create an APEX workspace:

1. In the APEX Workspaces page, at the top right, click **Create Workspace**.
2. Enter the following fields:
 - **Workspace name:** Enter a name for the workspace.
 - **Database user:** Select a schema from the drop-down list, or type to enter a new schema.
 - **Workspace password:** Enter a password for the workspace.
 - **APEX Administrator:** Enter the administrator user details.

Show code: Select this option to view the PL/SQL code equivalent of the Create Workspace slider. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Workspace slider.

3. Click **Create**.

The new workspace is displayed in the APEX workspaces page. The workspace card includes details such as the number of applications in the workspace, the number of developers using the workspace, and the number of administrators responsible for the workspace. Depending on your role, the  **Actions** context menu provides options such as creating and deleting a workspace, and viewing schema details.

12

The Data Pump Page

Note

This feature is only available for Oracle Database 12.2 and later releases.

The Data Pump page enables you to monitor Data Pump jobs that were initiated through the available Database API endpoints, the DBMS_DATAPUMP package, or the SQL Developer Data Pump Export and Import wizards.

To import data using Data Pump, click **Import Data**. See [Importing Data](#).


The section at the top displays the total number of executing jobs, stopped jobs, and completed jobs. Click a tile (example, STOPPED) to filter and view the corresponding list of STOPPED jobs in the default card format.

You can filter or sort the jobs and set the time period by which to refresh the data.

A job card displays the following details: Job name, import or export operation, percentage of completion and time elapsed, and links to dump files and logs. The status of the job is indicated by the colour of the icon on the left side of the card. Green indicates successful jobs, yellow indicates that the jobs need to be reviewed, and blue indicates that the jobs are in progress.

In a job card, you can:

- Use  **Download** to access dump files for completed jobs.
- Use  **Log** to access the log files.



Data Pump

1 EXECUTING 1 STOPPED 1 COMPLETED

Search by Job Name, Owner Filter by Sort by Refresh 10 sec...

DATAPUMP_REST_EXPORT_202108101...
EXPORT
35% 12:30:42.000 elapsed
POBDBA
EXPDAT%u-20210810_23_28_17.DMP
http://localhost:8080/ords/adminUser/_/db-api/stable/database/datapump/jobs/P...

DATAPUMP_REST_EXPORT_202108110...
EXPORT
50% 00:27:09.000 elapsed
POBDBA
EXPDAT%u-20210811_14_18_17.DMP
http://localhost:8080/ords/adminUser/_/db-api/stable/database/datapump/jobs/P...

DATAPUMP_REST_EXPORT_202108110...
EXPORT
100% completed in 00:00:35.000
POBDBA
EXPDAT%u-20210811_14_31_03.DMP
http://localhost:8080/ords/adminUser/_/db-api/stable/database/datapump/jobs/P...

12.1 Importing Data

This section provides the steps for importing data using Oracle Data Pump into your on-premises database.

Note

Before you begin, you must have an export job. To create an export job using cURL, see [Create an Export Data Pump Job](#) in *REST APIs for Oracle Database*.

1. In the Data Pump page, on the top right, click **Import**.
The Import wizard appears.
2. In the Source step, enter the following fields:
 - a. **Directory:** Select the directory that contains the source dump files.
 - b. **Import Pattern:** Type the import pattern.Click **Next**.
3. In the Import step, enter the following fields:
 - **Import Name:** Enter a name for the import job.
 - **Import Type:** Select the type of import. The options are Full, Tables, Schemas, and Tablespaces.

Note

If you select **Full**, you skip the Filter step in the wizard and directly go to the Mapping step.

- **Content:** Select **Data Only**, **DDL Only**, or **Data and DDL**.
 - **Encrypt:** Select if encrypted and enter an encryption password.
- Click
- Next**
- .
4. In the Filter step, depending on the import type, all the schemas, tables, or tablespaces for the import job are listed. Select the ones that apply. Click **Next**.
 5. In the Mapping step, select the source schema and enter a new name for the target schema. If needed, do the same for tablespaces. Click **Next**.
 6. In the Options step, enter the following fields:
 - **Threads:** Specify the maximum number of threads of active execution operating on behalf of the import job. The default is 1.
 - **Action on Table if Table Exists:** Specify the action needed if that table that import is trying to create already exists.
 - **Skip Unusable indexes:** Select to specify whether the import skips loading tables that have indexes that were set to the Index Unusable state.
 - **Regenerate Object IDs:** Select to create new object identifies for the imported database objects.

- **Delete Master Table:** Select to indicate whether the Data Pump control job table should be deleted or retained at the end of an Oracle Data Pump job that completes successfully.
- **Overwrite Existing Datafiles:** Select to indicate that if a table already exists in the destination schema, overwrite it.
- **Version:** Select the version of database objects to import.
- **Logging:** Select to create a log file. Enter the log directory and log file name.

Import

✓

✓

3

✓

5

6

Source

Import

Filter

Mapping

Options

Summary

Threads

1

Action on Table if Table Exists

☒ None
☐ Skip
☐ Append
☐ Truncate
☐ Replace

Skip Unusable Indexes

☐

Regenerate Object IDs

☐

Overwrite Existing Datafiles

☐

Delete Master Table

☐

Append Timestamp to Log and Job Names

☒

Version

Compatible

Logging

☒

Log Directory

DATA_PUMP_DIR

Log File Name

IMPORT

Click **Next**.

7. The Summary step displays a summary of all the selections made in the previous steps.

Select **Show Code** at the bottom to see the PL/SQL code equivalent of the form.

Import

✓

✓

3

✓

✓

6

Source
Import
Filter
Mapping
Options
Summary

```

DECLARE
  L_JOB_STATE VARCHAR2(1000);
  L_JOB_HANDLE number;
  L_HAS_FAILED VARCHAR2(100):= 'ERROR';
  L_SHOULD_TRY_TO_GET_STATUS number := 0;
  L_TABLE_EXISTS NUMBER;
BEGIN
  L_JOB_HANDLE := dbms_datapump.open( operation => 'IMPORT', job_mode => 'FULL', job_name => 'ON_PREM_TEST_17_11_51', version =>
  L_SHOULD_TRY_TO_GET_STATUS := 1;
  dbms_datapump.set_parallel( handle => L_JOB_HANDLE, degree => 1);
  dbms_datapump.add_file( handle => L_JOB_HANDLE, filename => 'IMPORT_17_11_51.LOG', directory => 'DATA_PUMP_DIR', filetype => D
  dbms_datapump.set_parameter( handle => L_JOB_HANDLE, name => 'KEEP_MASTER', value => 1 );
  dbms_datapump.add_file( handle => L_JOB_HANDLE, filename => 'EXPDAT%u-14_30_07.DMP', directory => 'DATA_PUMP_DIR', filetype =>
  dbms_datapump.set_parameter( handle => L_JOB_HANDLE, name => 'INCLUDE_METADATA', value => 1 );
  dbms_datapump.set_parameter( handle => L_JOB_HANDLE, name => 'DATA_ACCESS_METHOD', value => 'AUTOMATIC' );
  dbms_datapump.set_parameter( handle => L_JOB_HANDLE, name => 'REUSE_DATAFILES', value => 0);
  dbms_datapump.set_parameter( handle => L_JOB_HANDLE, name => 'SKIP_UNUSABLE_INDEXES', value => 0 );
  dbms_datapump.start_job( handle => L_JOB_HANDLE, skip_current => 0, abort_step => 0 );
  dbms_datapump.wait_for_job( handle => L_JOB_HANDLE, job_state => L_JOB_STATE );
  dbms_datapump.detach( handle => L_JOB_HANDLE );
  L_HAS_FAILED := 'NO_ERROR';
EXCEPTION
  WHEN OTHERS THEN
    IF ( ( L_HAS_FAILED = 'ERROR' ) AND ( L_SHOULD_TRY_TO_GET_STATUS = 1 ) ) THEN
      DBMS_DATAPUMP.DETACH( L_JOB_HANDLE );
    END IF;

    EXECUTE IMMEDIATE
      'SELECT COUNT(*) FROM DBA_TABLES WHERE TABLE_NAME=' ||
      DBMS_ASSERT.ENQUOTE_LITERAL ( 'ON_PREM_TEST_17_11_51' )
      INTO L_TABLE_EXISTS;

    IF L_TABLE_EXISTS <> 0 THEN
      EXECUTE IMMEDIATE
        'DROP TABLE' || DBMS_ASSERT.ENQUOTE_NAME ( 'ON_PREM_TEST_17_11_51', FALSE );
    END IF;
END;

```

?

☒
Show code

Back
Next
Import
Cancel

Click **Import**.

The start of the job execution is displayed on the Data Pump page.

Part IV

Monitoring

This part provides information about the monitoring features available with Database Actions.

Topics:

- [The Monitoring Pages](#)

The Monitoring Pages

The Monitoring menu provides access to several pages to view the performance and other characteristics of your database.

- [The Monitoring Overview Page](#)
- [The Performance Hub Page](#)
- [The Instance Viewer Page](#)
- [The Logins Page](#)
- [The Alerts Page](#)
- [The Sessions Page](#)
- [The Storage Page](#)
- [The Parameters Page](#)
- [The Real Time SQL Monitoring Page](#)
- [The Top SQL Page](#)
- [The Waits Page](#)
- [The AWR Page](#)

13.1 The Monitoring Overview Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Monitoring Overview page displays general information about the database.

To navigate to the Overview page, click **Selector**  to see the navigation menu and then select **Monitoring**.

The widgets on this page show snapshot information about the database status, online database storage, sessions, wait events, user accounts, alerts, and expiring passwords. Click a widget to go to its page where you can see a more detailed view of the data.

- **Database Status:** Shows the status of the database.
- **Used Online Database Storage:** Displays how much storage is being used by the database. You can click the title to open [The Storage Page](#).
- **Sessions:** Displays the number of sessions by session status. You can click the title to open [The Sessions Page](#).
- **Waits:** Displays how many wait events are occurring in the database by wait event class. When you click the title:
 - For Oracle Database 19c and later releases, the Performance Hub page is displayed.

- For Oracle Database 18c and previous releases, the Activity-Waits page is displayed.
- **Users:** Displays how many user accounts are in the open, locked and expired statuses. You can hover over a status to see a list of the user accounts with that status.
 - **Open:** This status indicates that the user's account is unlocked and access to the database is enabled.
 - **Locked:** This status indicates that the user's account is locked and access to the database is disabled. The account must be unlocked to enable access to the database.
 - **Expired:** This status indicates that the user's password has expired and must be changed before the user can log in to the database.
- **Alerts:** Displays a summary of alerts over the last 7 days. You can click the title to open [The Alerts Page](#)
- **Expiring Passwords:** Shows a list of user accounts and whether a user account password has expired or the number of days before it will expire.

13.2 The Performance Hub Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles for Oracle Database 19c and later releases.

The Performance Hub page shows performance data for a time period you specify.

To navigate to the Performance Hub page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Performance Hub**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Performance Hub**.

Note

The Performance Hub page is available in the following user interface languages: French, Japanese, Korean, Traditional Chinese, and Simplified Chinese. If you change the language to German, Spanish, Italian, or Portuguese in Preferences, the Performance Hub page reverts to English.

The Performance Hub page consists of these parts:

- **Time Range Area:** Use the controls in time range area at the top of the page to specify the time period for which you want to view performance data.
- **ASH Analytics Tab:** Use this tab to explore ASH (Active Session History) information across a variety of different dimensions for the specified time period.
- **SQL Monitoring Tab:** Use this tab to view the top 100 SQL statement executions by different dimensions for the specified time period, and to view details of SQL statement executions you select.

See [Performance Hub](#) in the Database service documentation for more information.

13.3 The Instance Viewer Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.


The Instance Viewer is a performance dashboard that enables a DBA user to see a graphical representation of information about the instance associated with a connection.

To navigate to the Instance Viewer page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Instance Viewer**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Instance Viewer**.

When hovering over a chart, a dark blue border indicates that you can click to drill down to the relevant page for more detailed information.

Figure 13-1 Dark Blue Border to Drill Down



The types of information displayed include:

- Database
- Clients
- Sessions
- Processes (Counts, Execution Rate, Parse Rate, Open Cursors, Commit Rate)
- Waits
- Memory (DB Block Rate, Logical Reads, Allocation, Redo Generation)
- Storage (Files, Redo Log)
- DB CPU Ratio (database operations as a percentage of CPU activity)

- Top SQL (provides performance metrics on SQL operations, which are sorted based on the CPU time required)

13.4 The Logins Page

Note

Available only if the following three conditions are met:

- You are signed in as a database user with **DBA** and **PDB_DBA** roles for Oracle Database 12c and later releases.
- Auditing is enabled in the database.
- At least one of the four actions displayed on the Logins page must have an audit policy that is enabled. By default, Failed Logins is enabled.

The Logins page shows the number of successful logins, failed logins, timed-out logoffs, and logoffs that have occurred within the last hour.


To navigate to the Logins page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Logins**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Logins**.

The Last Hour Statistics widget only displays the actions being audited. If any of the tabs (successful logins, failed logins, timedout logoffs or logoffs) are disabled, the database is currently not auditing that action. If you want to activate the action, you must create an audit policy with the desired action. To create an audit policy, see Auditing Activities.

Click a tab in Last Hour Statistics to view the corresponding entries in the table.

Figure 13-2 Tabs in Last Hour Statistics



You can search for results in the table by entering values in the Column, Operator and Value fields and clicking Filter . You can also filter by selecting one or more of the available filter options for each action.

Click Refresh  to refresh the data in the grid.

The following figure shows the three filters for the Failed Logins action.

Figure 13-3 Filters for Failed Logins

Last Hour Statistics

SUCCESSFUL LOGINS	0 FAILED LOGINS
-------------------	--------------------

Column	Operator	Value
OS User	equals	
☒ action_name equals 'LOGON'	⊗	
☒ returncode doesn't equal '0'	⊗	
☒ time greater than '12/9/2019, 8:41:54 AM'	⊗	

OS User	User	Date
---------	------	------

To return to the previous page, remove filters by clicking Remove Filter  or refresh the browser.

13.5 The Alerts Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Alerts page is a chronological log of messages and errors and is commonly used to learn whether the background processes have encountered errors. You can review the alert log periodically to verify that your database system is operating normally.

To navigate to the Alerts page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Alerts**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Alerts**.

The alert log includes the following:

- Nondefault initialization parameters used at startup
- Administrative operations, such as STARTUP, SHUTDOWN, ARCHIVE LOG, RECOVER, and CREATE/ALTER/ DROP DATABASE/TABLESPACE
- Messages and errors relating to the functions of certain background processes, such as LGWR
- Internal errors (ORA-600), block corruption errors (ORA-1578), and deadlock errors (ORA-60)

Click Refresh  at the top right of the page to refresh the data.

You can search for a specific value in the log by selecting the display column in the first drop-down list, selecting the condition in the second drop-down list, entering the search value in the box, and clicking the search icon.

In the display table, if you right-click the header row, you see:

Columns: Enables you to select columns to show or hide.

Sort: Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

If you right-click any other part of the display table, you see:

Count Rows: Displays the number of rows in the table.

Single Record View: Enables you to view data for a table or view, one record at a time.

Copy: Enables you to copy data from a cell or a row or a range of rows. To copy from more than one row, select the rows you want to copy, right-click by pressing the SHIFT or CTRL key, and select **Copy**.

13.6 The Sessions Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Sessions page shows information about all currently open sessions in the database.

To navigate to the Sessions page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Sessions**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Sessions**.

The data is automatically refreshed at intervals ranging from 10 seconds to 2 minutes. You can also refresh the data by clicking Refresh  at the top right of the screen.

The table shows summarized data about each open session. Select a session in the table to see more detailed data in the Session Details table below, such as the last SQL statement, explain plan, waits, contention, and so on. You can use the Column, Operator and Value fields to search for the required sessions.

In the display table, if you right-click the header row, you see:

Columns: Enables you to select columns to show or hide.

Sort: Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

If you right-click any other part of the display table, you see:

Count Rows: Displays the number of rows in the table.

Single Record View: Enables you to view data for a table or view, one record at a time.

Copy: Enables you to copy data from a cell or a row or a range of rows. To copy from more than one row, select the rows you want to copy by pressing the SHIFT or CTRL key, right-click and select **Copy**.

13.7 The Storage Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Storage page shows the storage used based on the current allocation of tablespaces along with additional drill-down capabilities to view segments.

To navigate to the Storage page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Storage**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Storage**.

You can refine the list of segments shown by using the filter feature. Click **View Datafiles** to view the datafiles in each tablespace.

You can view tablespace and segment space usage.

To view space usage information

1. From the Database drop-down menu, click **Storage**.

The Storage page displays. If the Oracle database is version 12c or later, the Storage page shows the used and allocated storage space for tablespaces in any pluggable database. If the Oracle database is version 11g, the Storage page shows the used and allocated space for the entire database.

2. You can click a tablespace to view its storage information. An interactive report appears, showing the segments that exist within the tablespace. Most segments are user objects, and they include tables, LOBs, and indexes.
3. On the Segments page, you can refine the list of segments shown by using the filter feature.

For example, you can search for all the segments for a specific owner (schema) by selecting OWNER from the first drop-down list, entering the owner (schema) name in the box, and clicking the search icon.

13.8 The Parameters Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Parameters pages displays initialization parameters, which are used to configure the database instance, including memory structures, and define locations for database files.

To navigate to the Parameters page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Parameters**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Parameters**.

Values for initialization parameters are stored in a text-based initialization parameter file (PFILE) or binary server parameter file (SPFILE). The initialization parameter file is read at database instance startup.

Click Refresh  at the top right of the page to refresh the data.

To perform a search, enter values in the search criteria columns and click the search icon to locate the initialization parameter.

In the display table, if you right-click the header row, you see:

Columns: Enables you to select columns to show or hide.

Sort: Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

If you right-click any other part of the display table, you see:

Count Rows: Displays the number of rows in the table.

Single Record View: Enables you to view data for a table or view, one record at a time.

Copy: Enables you to copy data from a cell or a row or a range of rows. To copy data from more than one row, select the rows you want to copy by pressing Shift or Ctrl, right-click and select **Copy**.

13.9 The Real Time SQL Monitoring Page

Note

Available for database users with DBA and PDB_DBA roles for Oracle Database 18c and previous releases. For non-administrator users, this feature is available only for Oracle Database 19c and later releases.

The Real Time SQL Monitoring page shows in real time the SQL statements that are being monitored in the database.

To navigate to the Real Time SQL Monitoring page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Real Time SQL Monitor**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Real Time SQL Monitor**.

In the Auto Refresh drop-down list, you can select the time (in seconds) to periodically refresh the data. Select 0 seconds to disable the auto-refresh.

This tool helps identify run-time issues for SQL statements by providing two major functions:

- General view of monitored statements
- View of SQL execution details

General View of Monitored Statements

The page contains a table of SQL statements currently running. This table shows the following information:

- **Status:** Current state of the SQL statement execution. For example, a SQL statement that has already finished its execution will show a status of "DONE".
- **Duration:** Amount of time a SQL statement is taking, or has taken, to execute.
- **SQL ID:** SQL identifier of the statement being monitored.
- **Session ID:** Session identifier that is executing, or has executed, the SQL statement.
- **Session Serial Number:** Uniquely identifies a session's objects.
- **Instance Degree of Parallelism:** This Degree of Parallelism (DOP) column shows how many instances and parallel execution servers are allocated. It is shown in the form of "number of instances" | "number of parallel servers".
- **Database Time:** Place the cursor over the database time to see a breakdown of the time and wait events.
- **CPU Time:** CPU time consumed by the execution of the query.
- **I/O Time:** I/O time consumed by the execution of the query.
- **Start Time:** Time in which the execution of the SQL statement started.
- **SQL Statement:** SQL statement being monitored.

Figure 13-4 Real Time SQL Monitoring Table


	STATUS	DURATION	TYPE	SQL ID	PLAN HASH	USERNAME	PARALLEL	DATABASE TIME	IO REI
1	✓ DONE	21 s	SQL	g6px76dmjv1jy	3702721588			19.81 s	
2	✓ DONE	20 s	SQL	g6px76dmjv1jy	3702721588			19.85 s	
3	✓ DONE	20 s	SQL	g6px76dmjv1jy	3702721588			19.7 s	
4	✓ DONE	20 s	SQL	g6px76dmjv1jy	3702721588			19.35 s	
5	✓ DONE	20 s	SQL	g6px76dmjv1jy	3702721588			19.57 s	
6	✓ DONE	20 s	SQL	g6px76dmjv1jy	3702721588			19.54 s	
7	✓ DONE	21 s	SQL	g6px76dmjv1jy	3702721588			20.09 s	
8	✓ DONE	22 s	SQL	g6px76dmjv1jy	3702721588			19.96 s	
9	✓ DONE	22 s	SQL	g6px76dmjv1jy	3702721588			19.23 s	
10	✓ DONE	21 s	SQL	g6px76dmjv1jy	3702721588			19.46 s	
11	✓ DONE	21 s	SQL	g6px76dmjv1jy	3702721588			18.96 s	
12	✓ DONE	21 s	SQL	g6px76dmjv1jy	3702721588			19.22 s	

View of SQL Execution Details

When a SQL statement is drilled down from the main monitor table, a detailed view is shown. The SQL ID, Start Time and the SQL Execution ID represent the execution key that uniquely identify this SQL statement. A detail view consists of the general characteristics that integrate the execution of a SQL statement.

General information about the query execution is provided:

- **Execution Plan:** Degree of Parallelism of the SQL statement
- **Execution Started:** Time that the SQL statement execution started
- **Last Refresh Time:** Last update time of the SQL monitor registry for the SQL statement
- **Execution ID:** Execution identifier
- **User:** User in the format USER@CONTAINER
- **SQL Text:** Formatted view of the SQL statement that is being executed.

Figure 13-5 SQL Execution Details in Real Time SQL Monitor

General statistics of the SQL statement are provided: total duration of execution, the number of buffered gets, number of Input/Output requests and bytes.

Detailed information of the statement: This space holds the information corresponding to the explain plan, parallel behavior and CPU activity involved in the execution of the statement:

- **Plan Statistics:** Explain plan of the execution of the SQL statement in the form of a table. Each row is a different operation involved in the execution of the SQL statement and it shows hierarchy dependency by adding a space at the beginning of the text in the Operation column.
- **Parallelism Details for the SQL statement:** Each execution consists of a parallel coordinator and one or more parallel sets. Each set can have one or more processes. When a row has dependents, each of its columns will be the sum of the values of its dependants. When this happens, a sigma symbol will appear to show that a value consists of the sum of others.

Note

For more information, see "Monitoring the Database" in the *Oracle Database Administrator's Guide*.

13.10 The Top SQL Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The Top SQL page displays SQL statements executed in the database based on CPU time consumed. The SQL statement that consumes the maximum CPU time is right at the top, and the remaining statements continue in descending order of CPU time.

To navigate to the Top SQL page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Top SQL**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Top SQL**.

This feature enables you to focus your SQL tuning efforts on the statements that can have the most impact on database performance.

Click a row to see a formatted view of the SQL statement, the execution plan, the runtime history from Active Session History, and a SQL Tuning Advisor report. Active Session History is part of the Diagnostics Pack, and SQL Tuning Advisor is part of the Tuning Pack for Oracle Database.

13.11 The Waits Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles for Oracle Database 18c and previous releases.


The Waits page shows a chart with the distribution of wait events in real time.

To navigate to the Waits page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Waits**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **Waits**.

Use the slider controls to zoom in on a specific time period. To use the slider controls, place the cursor over the handles at both sides of the box and drag the sides to the time period required. The chart above will refresh to the selected time period. The table will also automatically refresh and the wait events will filter to that period of time enabling you to easily identify the problem SQL statement.

Figure 13-6 Distribution of Wait Events Chart



When you place the cursor over data points in the chart, a pop-up box displays details about the wait event.

Figure 13-7 Display Wait Event Details



13.12 The AWR Page

Note

Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles.

The AWR page in Database Actions enables you to generate, view and download Automatic Workload Repository (AWR) reports.

To navigate to the AWR page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **AWR**.
- Click **Selector**  to display the navigation menu. Under Monitoring, select **AWR**.

The Automatic Workload Repository collects, processes, and maintains performance statistics for the database. The gathered data can be displayed in reports and views.

To generate or view a report in the AWR page, select the range of the time period required using the **Start ID** and **End ID** fields. The snapshot dropdown list for each field is sorted in descending order and starts with the most recent database snapshot.

Note

The **Generate Report** button is available only when the value in the Start ID field is lower than the value in the End ID field.

Monitoring > AWR Report

Start ID

3

End ID

4

Generate Report

Download Report

WORKLOAD REPOSITORY PDB report (root snapshots)

DB Name	DB Id	Unique Name	Role	Edition	Release	RAC	CDB
DB213C	285474074	DB213C	PRIMARY	EE	21.0.0.0.0	NO	YES
Instance	Inst Num	Startup Time	User Name	System Data Visible			
DB213C	1	06-Jun-22 06:09	PDBDBA	YES			
Container DB Id		Container Name			Open Time		
345142159		DB213P			06-Jun-22 06:09		
Host Name	Platform	CPU	Cores	Sockets	Memory (GB)		
db213	Linux x86 64-bit	48	24	2	500		
	Snap Id	Snap Time		Sessions	Cursors/Session		
Begin Snap:	3	06-Jun-22 07:58:33		36			
End Snap:	4	06-Jun-22 08:58:47		29			

A

Supported SQL*Plus and SQLcl Commands in SQL Worksheet

The following sections list the SQL*Plus and SQLcl commands supported in the worksheet.

A.1 Supported SQL*Plus Commands

The SQL worksheet supports most of the SQL*Plus commands except those statements that are related to formatting.

- / (slash)
- @@ { url | file_name[.ext] } [arg ...]
- ACC[EPT] variable [NUM[BER] | CHAR | DATE | BINARY_FLOAT | BINARY_DOUBLE] [FOR[MAT] format] [DEF[AULT] default] [PROMPT text | NOPR[OMPT]] [HIDE]
- ARCHIVE LOG LIST
- BRE[AK] [ON report_element [action [action]]] ...
- BTI[TLE] [printspec [text | variable] ...] | [ON | OFF]
- C[HANGE] sepchar old [sepchar [new [sepchar]]]
- CL[EAR] option ...
- COL[UMN] [{column | expr} [option ...]]
- COMP[UTE] [function [LAB[EL] text] ... OF {expr | column | alias} ...ON {expr | column | alias | REPORT | ROW} ...]
- COPY {FROM database | TO database | FROM database TO database} {APPEND | CREATE | INSERT | REPLACE} destination_table[(column, column, column, ...)] USING query
- DEF[INE] [variable] | [variable = text]
- DESC[RIBE] {[schema.]object[@connect_identifier]}
- DISC[ONNECT]
- EXEC[UTE] statement
- {EXIT | QUIT} [SUCCESS | FAILURE | WARNING | n | variable | :BindVariable] [COMMIT | ROLLBACK]
- GET [FILE] file_name[.ext] [LIST | NOLIST]
- HO[ST] [command]
- L[IST] [n | n m | n * | n LAST | * | * n | * LAST | LAST]
- PAU[SE] [text]
- PRINT [variable ...]
- PRO[MPT] [text]

- {QUIT | EXIT} [SUCCESS | FAILURE | WARNING | n | variable | :BindVariable] [COMMIT | ROLLBACK]
- R[UN]
- SAV[E] [FILE] file_name[.ext] [CRE[ATE] | REP[LACE] | APP[END]]
- SET system_variable value
- SHO[W] [option]
- SHUTDOWN [ABORT | IMMEDIATE | NORMAL | TRANSACTIONAL [LOCAL]]
- STA[RT] { url | file_name[.ext] } [arg ...]
- STARTUP db_options | cdb_options | upgrade_options
- TIMI[NG] [START text | SHOW | STOP]
- TTI[TLE] [printspec [text | variable] ...] | [ON | OFF]
- UNDEF[INE] variable ...
- VAR[IABLE] [variable [type][=value]]
- XQUERY xquery_statement

A.2 Supported SQLcl Commands

The SQL worksheet supports many of the SQLcl commands.

- ALIAS
- APEX
- BRIDGE
- CTAS
- DDL
- FORMAT
- INFORMATION
- LOAD
- NET
- OERR
- RESERVED_WORDS
- SCRIPT
- SETERRORL
- SODA (See [SODA Commands](#))
- TNSPING
- TOSUB
- WHICH

A.2.1 SODA Commands

SODA (Simple Oracle Document Access) commands are supported in the SQL code editor. SODA allows schemaless application development using the JSON data model. The commands are:

- SODA create *<collection_name>* — Creates a new collection
- SODA list — Lists all the collections
- SODA get *<collection_name>* [-all | -f | -k | -klist] [{*<key>* | *<k1>* *<k2>* ... | *<qbe>*}] — Lists documents in the collection. Optional arguments:
 - all: Lists the keys of all documents in the collection
 - k: Lists documents matching the specific *<key>*
 - klist: Lists documents matching the list of keys
 - f: Lists documents matching the *<qbe>*
- SODA insert *<collection_name>* *<json_str | filename>* — Inserts a new document within a collection
- SODA drop *<collection_name>* — Deletes existing collection
- SODA count *<collection_name>* [*<qbe>*] — Counts number of documents inside collection. Optional parameter *<qbe>* returns number of matching documents
- SODA replace *<collection_name>* *<oldkey>* *<new_{str | doc}>* — Replaces one document with another
- SODA remove *<collection_name>* [-k | -klist | -f] [{*<key>* | *<k1>* *<k2>* ... | *<qbe>*}] — Removes documents from collection. Optional arguments:
 - k: Removes document in collection matching the specific *<key>*
 - klist: Removes document in collection matching the list *<key1>* *<key2>* ...
 - f: Removes document in collection matching *<qbe>*