

Using Oracle® SQL Developer Web



Release 20.3

F34127-01

October 2020

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Using Oracle® SQL Developer Web, Release 20.3

F34127-01

Copyright © 2018, 2020, Oracle and/or its affiliates.

Primary Author: Celin Cherian

Contributing Authors: Mike Fitch, Chuck Murray

Contributors: Jeff Smith, Syme Kutz, Victor Martinez, Julio Tinajero, Renato Uriel Gutierrez Salas, Antonio Hernandez, Sofia Isadora Padilla Munoz, Miguel Sanchez Villafan, Viktor Perelmouter, Tulika Das, Eric Paapanen

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 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" or "commercial computer software documentation" 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 and Java 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.

Contents

Preface

Audience	vii
Documentation Accessibility	vii
Product Accessibility	vii
Related Documents	vii
Conventions	viii
Third-Party License Information	viii

1 Changes in Release 20.3 for Using Oracle SQL Developer Web

2 About Oracle SQL Developer Web

2.1 About the SQL Developer Web User Interface	2-1
2.2 Accessing Oracle SQL Developer Web	2-3
2.3 The Database Actions Home Page	2-4

3 The SQL and Data Modeler Pages

3.1 The Overview Page	3-1
3.2 The SQL Page	3-2
3.2.1 Navigating Objects and Worksheets	3-2
3.2.2 Executing SQL Statements in the SQL Editor	3-5
3.2.2.1 Keyboard Shortcuts	3-7
3.2.3 Viewing the SQL Output	3-7
3.2.4 Loading Data	3-10
3.2.4.1 Loading Data from a Local File to an Existing Table	3-10
3.2.4.2 Loading Data from a Local File to a New Table	3-12
3.2.4.3 Format Specifications for JSON, AVRO, and XML Files	3-14
3.3 The Data Modeler Page	3-19
3.3.1 Navigating Diagrams and Objects	3-20
3.3.2 About the Data Modeler Editor	3-22
3.3.2.1 About Viewing JSON Data Guides	3-24

3.4	Creating and Editing Database Objects	3-28
3.4.1	The Table Properties Dialog	3-30
3.4.1.1	Columns Pane	3-32
3.4.1.2	Primary Key Pane	3-34
3.4.1.3	Unique Keys Pane	3-34
3.4.1.4	Indexes Pane	3-35
3.4.1.5	Foreign Keys Pane	3-36
3.4.1.6	Table Constraints Pane	3-36
3.4.1.7	Comments Pane	3-36
3.4.1.8	Storage Pane	3-37
3.4.1.9	External Table Properties Pane	3-37
3.4.1.10	Materialized View Pane	3-40
3.4.1.11	DDL Pane	3-42
3.4.1.12	Output Pane	3-42
3.4.2	The Index Properties Dialog	3-42
3.4.3	The Sequence Properties Dialog	3-43
3.4.4	The View Properties Dialog	3-44
3.4.5	The Synonym Properties Dialog	3-45
3.4.6	Implied Foreign Keys	3-46
3.4.7	The Materialized View Log Properties Dialog	3-49

4 The REST Pages

4.1	About RESTful Web Services	4-1
4.2	RESTful Services Terminology	4-1
4.3	About the Overview Page	4-2
4.4	About the REST Search Toolbar	4-3
4.5	Creating RESTful Web Services	4-4
4.5.1	Managing Resource Modules	4-4
4.5.1.1	Creating a Resource Module	4-5
4.5.1.2	Editing a Resource Module	4-6
4.5.1.3	Deleting a Resource Module	4-6
4.5.1.4	Publishing/Unpublishing a Resource Module	4-6
4.5.1.5	Exporting a Resource Module	4-6
4.5.2	Managing Resource Templates	4-7
4.5.2.1	Creating a Resource Template	4-7
4.5.2.2	Editing a Resource Template	4-8
4.5.2.3	Deleting a Resource Template	4-8
4.5.3	Managing Resource Handlers	4-8
4.5.3.1	Creating a Resource Handler	4-9
4.5.3.2	Example: Inserting a Record using a POST Handler	4-10

4.5.3.3	Editing a Resource Handler	4-12
4.5.3.4	Deleting a Resource Handler	4-13
4.5.4	Viewing Resource Handler Details and Managing Parameters	4-13
4.5.4.1	Creating a Parameter	4-14
4.5.4.2	Editing a Parameter	4-14
4.5.4.3	Deleting a Parameter	4-14
4.6	Securing RESTful Web Services	4-14
4.6.1	Managing Roles	4-15
4.6.1.1	Creating a Role	4-15
4.6.1.2	Editing a Role	4-15
4.6.1.3	Deleting a Role	4-16
4.6.1.4	Viewing Assigned Privileges	4-16
4.6.2	Managing Privileges	4-16
4.6.2.1	Creating a Privilege	4-16
4.6.2.2	Editing a Privilege	4-17
4.6.2.3	Deleting a Privilege	4-17
4.6.3	Managing OAuth Clients	4-17
4.6.3.1	Creating an OAuth Client	4-18
4.6.3.2	Example: Creating an OAuth Client	4-19
4.6.3.3	Editing an OAuth Client	4-20
4.6.3.4	Deleting an OAuth Client	4-21

5 The JSON Page

5.1	About the JSON User Interface	5-1
5.2	Managing JSON Collections	5-3
5.2.1	Creating a Collection	5-3
5.2.2	About Adding or Editing a JSON Document	5-5
5.3	About Querying Documents in a Collection	5-5
5.4	Managing JSON Search Indexes for a Collection	5-7
5.5	Viewing the JSON Data Guide Diagram for a Collection	5-8

6 The Database Users Page

6.1	About Create/Edit User	6-2
6.2	About REST Enable User	6-3
6.3	About Delete User	6-3

7 The Monitoring Pages

7.1	The Monitoring Overview Page	7-1
7.2	The Performance Hub Page	7-2

7.3	The Instance Viewer Page	7-4
7.4	The Logins Page	7-5
7.5	The Alerts Page	7-6
7.6	The Sessions Page	7-7
7.7	The Storage Page	7-8
7.8	The Parameters Page	7-9
7.9	The Real Time SQL Monitoring Page	7-9
7.10	The Top SQL Page	7-12
7.11	The Waits Page	7-12

A Supported SQL*Plus and SQLcl Commands

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 SQL Developer Web, a web-based interface of Oracle SQL Developer.

Audience

This online help is intended for those using Oracle SQL Developer Web included with Oracle REST Data Services.

Documentation Accessibility

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

Access to Oracle Support

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

Product Accessibility

Oracle SQL Developer Web provides features to support accessibility. See *SQL Developer Web Accessibility Guide*.

Related Documents

Oracle REST Data Services Installation, Configuration, and Development 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 Installation, Configuration, and Development Guide*

1

Changes in Release 20.3 for *Using Oracle SQL Developer Web*

This sections lists the changes in this release.

New Features

New features in this release include:

- Develop RESTful web services using the REST pages.
See [The REST Pages](#)
- Manage and query JSON collections using the JSON page.
See [The JSON Page](#)
- Display database object information in the SQL page.
See [Navigating Objects and Worksheets](#)

2

About Oracle SQL Developer Web

Oracle SQL Developer Web (also known as Database Actions) is a web-based application 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, monitoring databases, and creating Data Modeler diagrams.

Note:

Some features of Oracle SQL Developer Web 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.

SQL Developer Web is also available in Oracle Autonomous Database cloud services. For more information, see [Database Actions](#).

See Also:

- [About Create/Edit User](#) for creating users and assigning roles.
- [ORDS System Requirements](#) for system requirements.

2.1 About the SQL Developer Web User Interface

The SQL Developer Web 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

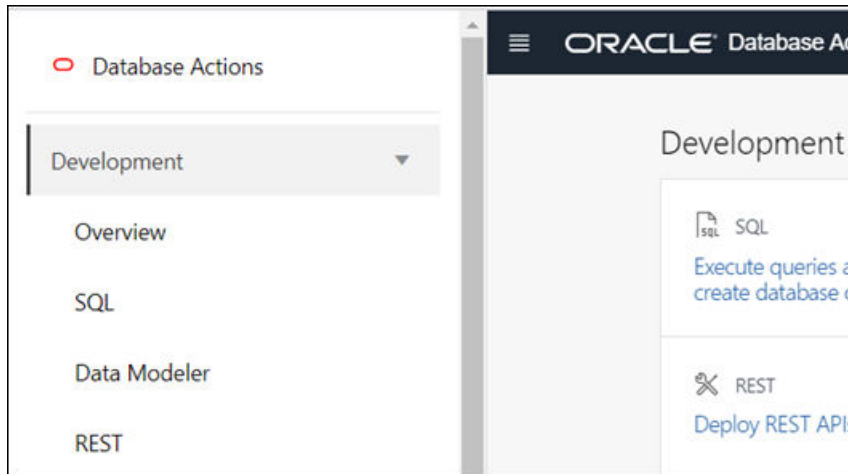
Header

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



- **Selector Icon**

Click the Selector icon  to see the main menu slide into view in a navigation pane. Use the Selector to switch between the Database Actions page and the navigation pane.



- **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:** Opens a dialog that lets you set your preferred time zone and user interface language. The languages available are English, German, Spanish, French, Italian, Japanese, Korean, Portuguese, and Chinese.
- **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.

2.2 Accessing Oracle SQL Developer Web

Oracle SQL Developer Web is included with Oracle REST Data Services.

To access SQL Developer Web:

1. Enable SQL Developer Web in Oracle REST Data Services. For more information, see *Advanced Installation Using Command-Line Prompts in Oracle REST Data Services Installation, Configuration, and Development Guide*.
2. To use SQL Developer Web, you must sign in as a database user whose schema has been REST-enabled for SQL Developer Web.

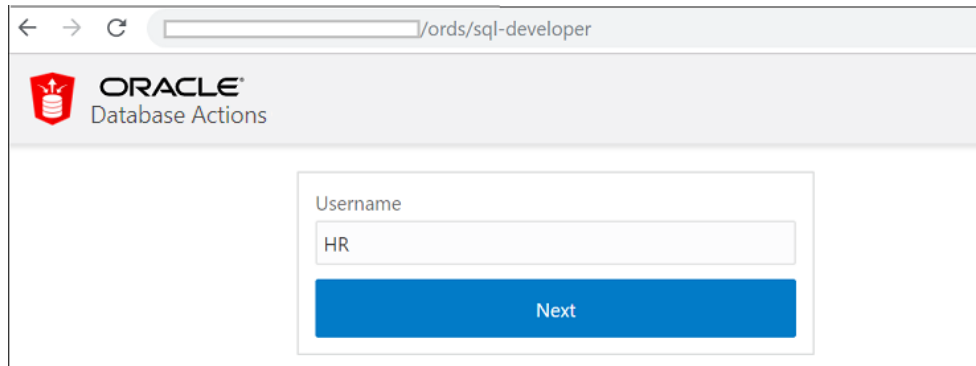
Execute the following code as a database user with the DBA role:

```
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 SQL Developer Web. 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. In the login page (`http://domain name/ords/sql-developer`), enter the name of the user and click **Next**.

Figure 2-1 Login Page



You see the URL change to `/ords/<Username>`. Enter the password, and click **Sign in**. The Database Actions page appears.

When you log out, you return to the login page (`http://domain name/ords/sql-developer`) and are prompted for the user name again.

2.3 The Database Actions Home Page

When you log in to SQL Developer Web, the home page is displayed containing links to the main features in the application.

The home page consists of three main groups: Development, Administration, and Monitoring. The Administration and Monitoring groups are available only for administrator user roles. Each group consists of feature-based icons that you can click to navigate to the respective pages.

The main features are:

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](#)
- **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](#)

Administration

- **Database Users:** Perform user management tasks such as create, edit, and REST enable users. Database Users is available only if you are signed in as a database user with DBA and PDB_DBA roles. See [The Database Users Page](#)

Monitoring

- Monitor database activity and performance using various tools. Monitoring is available only if you are signed in as a database user with DBA and PDB_DBA roles. See [The Monitoring 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** , click the down arrow to expand **Development** in the navigation pane, and then click **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 Database Actions page, click **SQL**.
- Click **Selector**  to display the left navigation pane. Expand the **Development** menu and 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:


- [Navigating Objects and Worksheets](#)
- [Executing SQL Statements in the SQL Editor](#)
 - [Keyboard Shortcuts](#)
- [Viewing the SQL Output](#)
- [Loading Data](#)

3.2.1 Navigating Objects and Worksheets

The Navigator tab in the left pane displays the saved objects for the selected schema.

[Figure 3-1](#) The following figure shows the various elements in the left pane.

Figure 3-1 Left Pane in SQL






- **Schema and Object Type selector:** Use the drop-down lists to select the schema and filter the results by object type.
- **Worksheets tab:** Displays the worksheets that are saved in the browser. You can open or delete an existing worksheet by using the context (right-click) menu.

 **Note:**

The worksheets created are dependent on the browser used (Internet Explorer, Mozilla Firefox, or Google Chrome) for the session at the time of saving the worksheet, and they will not be accessible from any other browser.

- **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 wildcard characters.
- **Context menu:** Options in the context menu are:
 - **Open:** Browse information relevant to the object type. For example, for a table, you can see the following: Columns, Dependencies, Constraints, Details, Grants, Indexes, Partition, Statistics and Triggers.

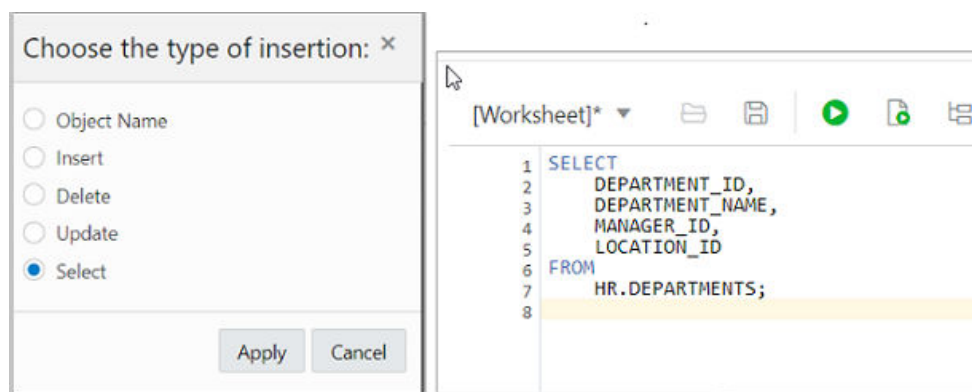
HR.COUNTRIES				
Columns				
Constraints		COLUMN_NAME	DATA_TYPE	NULLABLE
	1	COUNTRY_ID	CHAR(2 BYTE)	No
Dependencies	2	COUNTRY_NAME	VARCHAR2(40 BYTE)	Yes
Details	3	REGION_ID	NUMBER	Yes
Grants				
Indexes				
Partitions				
Statistics				
Triggers				

- **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.
- **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.

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 SQL Editor

The SQL editor enables you to enter SQL statements that you intend to execute.

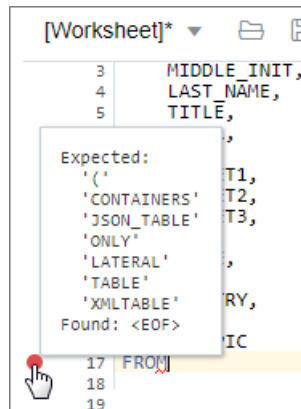
You can use SQL and PL/SQL statements to specify actions such as creating a table, inserting data, selecting data or deleting data from a table. For multiple statements, each non-PL/SQL statement must be terminated with either a semicolon or (on a new line) a slash (/), and each PL/SQL statement must be terminated with a slash (/) on a new line. SQL keywords are automatically highlighted.

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

If you press **Ctrl+Space**, the worksheet 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. Also, you can select multiple options in the list using **Ctrl+Click**.



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



The SQL toolbar contains icons for the following operations:

- **Worksheet** enables you to create a worksheet, open or delete an existing worksheet, and save a worksheet.

Note:

The worksheets are saved in the browser. Therefore, if you do sensitive work on a computer in a public network, remember to clear the browser cache before you leave. Additionally, the saved worksheets will only be displayed in the browser that was used at the time of creating the worksheet.

- **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 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.
- **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.
- **Autotrace** runs the statement and collects runtime 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.

For a list of shortcuts used in SQL, see [Keyboard Shortcuts](#).

3.2.2.1 Keyboard Shortcuts

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

Table 3-1 Keyboard Shortcuts for User Actions

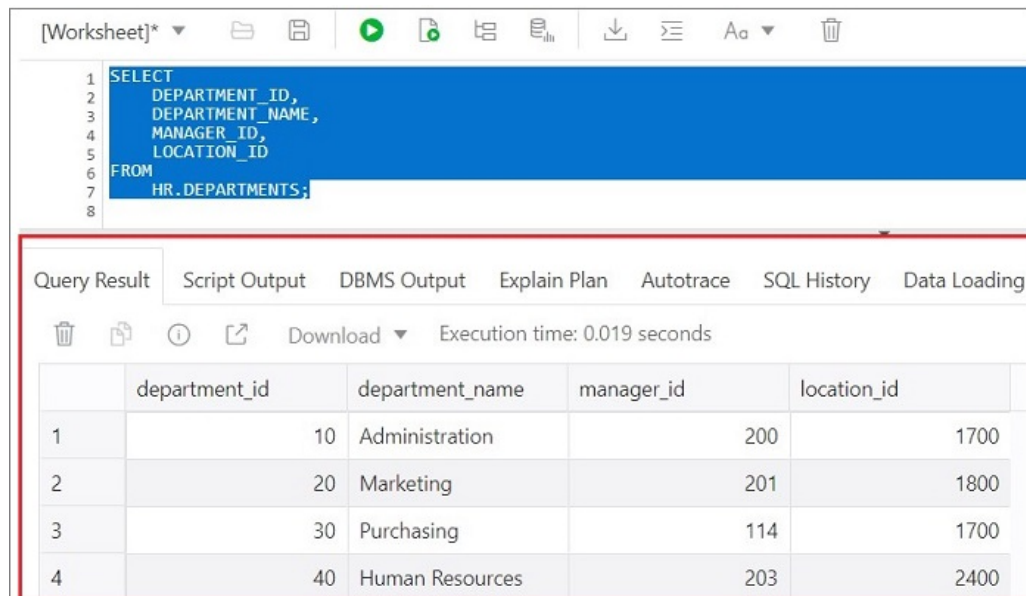
Keyboard Shortcut	Action
Ctrl+Enter/ Cmd+Enter	Runs the code as query.
Ctrl+Down Arrow/ Cmd+Down Arrow	Moves to the next SQL code from history.
Ctrl+Up Arrow/ Cmd+Up Arrow	Moves to the previous SQL code from history.
Ctrl+D/ Cmd+D	Clears the editor.
Ctrl+S/ Cmd+S	Saves the current worksheet.
Ctrl+O/ Cmd+O	Opens the worksheet browser dialog.
Ctrl+I/ Cmd+I	Downloads the content of the editor.
F1	Opens the help topic.
F5	Runs code as script.
F6	Shows Autotrace.
F10	Shows Explain Plan.
Ctrl+F7/ Cmd+F7	Formats code in the editor.
Ctrl+Space/ Cmd+Space	Autocompletes code (shows hints).
Windows+Esc/ Cmd+Esc	Focuses outside the editor and navigates to the rest of the application using the Tab key.

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 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.
- **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.
- **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.
- **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	OLS	(null)
OUTLN	OLNAME	(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

AR	Argentina
RU	
DE	
FR	
CA	
CH	Switzerland

3.2.4 Loading Data

In the SQL page, you can load data from a local file into an existing table or into a new table.

The file formats that you can load are CSV, XLS, XLSX, TSV, TXT, XML, JSON, and AVRO. For XML, JSON, and AVRO files, see [Format Specifications for JSON, AVRO, and XML Files](#).

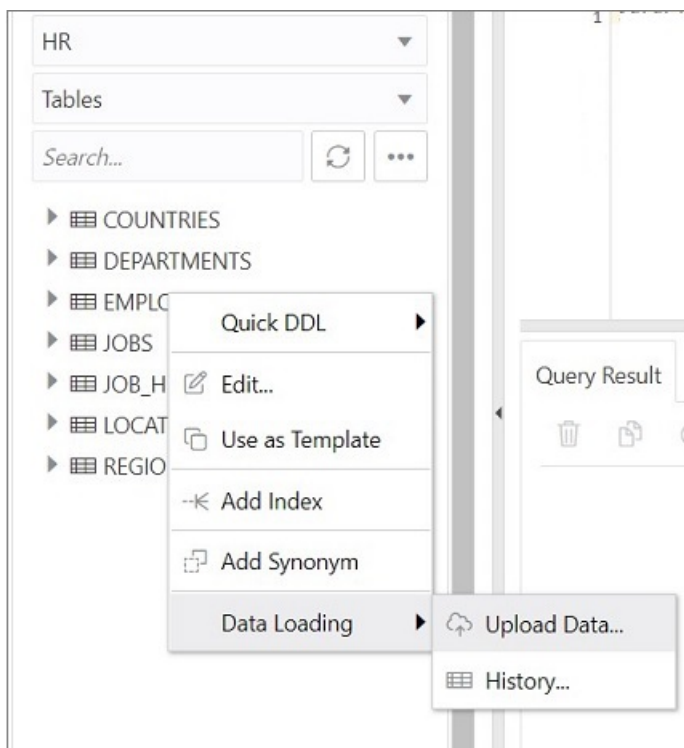
- [Loading Data from a Local File to a New Table](#)
- [Loading Data from a Local File to an Existing Table](#)

3.2.4.1 Loading Data from a Local File to an Existing Table

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

1. 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 3-6 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.

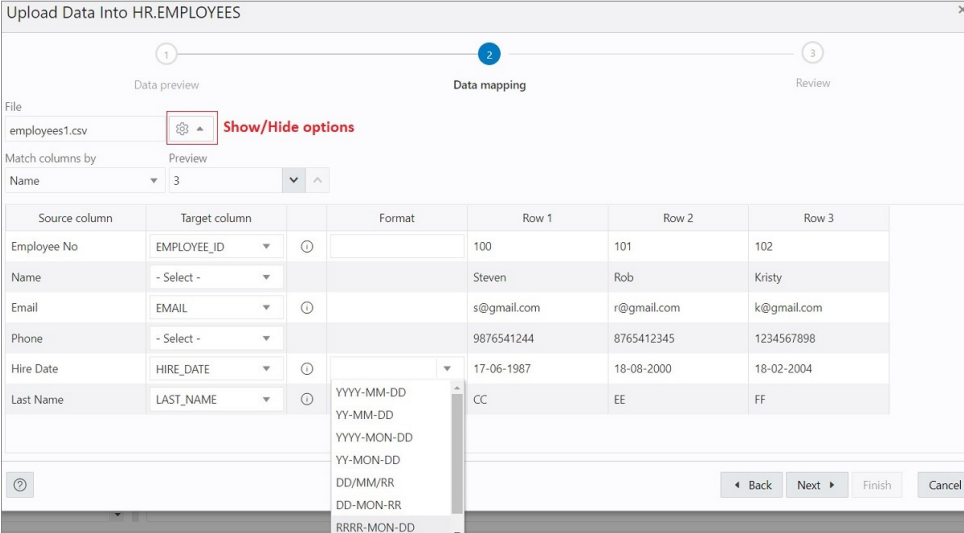
3. Click **Show/Hide options**  to display options that you can modify for data preview:
- **Column names:** Select **Get from file** to display column headers in the first row.
 - **Text enclosure** and **Field delimiter:** These options are visible only when the selected file is in plain text format (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 3-7 Data Mapping



Upload Data Into HR.EMPLOYEES

1 Data preview 2 Data mapping 3 Review

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

Match columns by: Preview

Name 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: 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 .

3.2.4.2 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 3-8 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`.

- In the worksheet output pane, select the **Data Loading** tab and drag and drop the local file into the output pane.

Figure 3-9 Data Loading Tab



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.
 - **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 3-10 Table Definition Step in Upload Data into New Table

Column Name	Column Type	Length/Precision	Scale	Default	Primary Key	Nullable	Format mask
<input checked="" type="checkbox"/> EMPLOYEE_NO	NUMBER				<input type="checkbox"/>	<input checked="" type="checkbox"/>	
<input checked="" type="checkbox"/> NAME	VARCHAR2	4000			<input type="checkbox"/>	<input checked="" type="checkbox"/>	
<input checked="" type="checkbox"/> EMAIL	VARCHAR2	4000			<input type="checkbox"/>	<input checked="" type="checkbox"/>	
<input type="checkbox"/> PHONE_	NUMBER				<input type="checkbox"/>	<input checked="" type="checkbox"/>	
<input checked="" type="checkbox"/> HIRE_DATE	TIMESTAMP WITH TIM...	4000			<input type="checkbox"/>	<input checked="" type="checkbox"/>	DD-MM-RRRR
<input checked="" type="checkbox"/> LAST_NAME	VARCHAR2	4000			<input type="checkbox"/>	<input checked="" type="checkbox"/>	


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 .

3.2.4.3 Format Specifications for JSON, AVRO, and XML Files

Data has to be stored in a particular format for JSON, AVRO, and XML files to load them successfully into a table.

The format specifications are described in the following sections.

3.2.4.3.1 JSON and AVRO Files

For JSON and AVRO files, the conversion for primitive types to table columns is supported only for top-level data. Nested objects are saved as JSON strings such as VARCHAR2 (JSON) or CLOB (JSON).

 **Note:**

JSON check constraints are available only for Oracle Database 12c and later releases.

Consider the following JSON file as an example:

```
[
  {
    "ItemNumber": 1,
    "Description": "One Magic Christmas",
    "Part": {
      "UnitPrice": 19.95,
      "UPCCode": 13131092899
    },
    "Quantity": 9,
    "Total": 179.55
  },
  {
    "ItemNumber": 2,
    "Description": "Lethal Weapon",
    "Part": {
      "UnitPrice": 17.95,
      "UPCCode": 85391628927
    },
    "Quantity": 5,
    "Total": 89.75
  }
]
```

The AVRO schema for this file:

```
{
  "type": "array",
  "items": {
    "type": "record",
    "fields": [
      {
        "name": "ItemNumber",
        "type": "int"
      },
      {
        "name": "Description",
        "type": "string"
      },
      {
        "name": "Part",
        "type": {
```

```

        "type": "record",
        "fields": [
            {
                "name": "UnitPrice",
                "type": "float"
            },
            {
                "name": "UPCCode",
                "type": "float"
            }
        ]
    },
    {
        "name": "Quantity",
        "type": "int"
    },
    {
        "name": "Total",
        "type": "float"
    }
]
}
}

```

Load the JSON file using "Upload Data" in the SQL page, and it is converted to the following table with two rows. `part` is a nested object that is assigned the column type CLOB (JSON) during data mapping.

	itemnumber	description	part	quantity	total
1	1	One Magic Christmas	{"UnitPrice":19.95,"UPCCode":131310...	9	179.55
2	2	Lethal Weapon	{"UnitPrice":17.95,"UPCCode":853916...	5	89.75

3.2.4.3.2 XML Files

This section lists the specifications for loading XML files.

- Attributes will have their own columns

If the XML data is structured as:

```

<?xml version="1.0"?>
<catalog>
  <book id="bk102">
    <author>Ralls, Kim</author>
    <title>Midnight Rain</title>
    <genre>Fantasy</genre>
    <publisher>John Doe</publisher>
  </book>
</catalog>

```

The generated columns in the table are `id`, `author`, `title`, `genre`, and `publisher`.

- Two or more levels of nesting are needed to parse the data
In the following example, the data that needs to be parsed will not be located because it has only one level of nesting (`catalog`).

```
<?xml version="1.0"?>
<catalog>
  <author>Ralls, Kim</author>
  <title>Midnight Rain</title>
  <genre>Fantasy</genre>
  <publisher>John Doe</publisher>
</catalog>
```

However, the following examples will work:

```
<?xml version="1.0"?>
<catalog>
  <book id="bk102">
    <author>Ralls, Kim</author>
    <title>Midnight Rain</title>
    <genre>Fantasy</genre>
    <publisher>John Doe</publisher>
  </book>
</catalog>
```

or

```
<?xml version="1.0"?>
<catalog>
  <bookstore>
    <book id="bk102">
      <author>Ralls, Kim</author>
      <title>Midnight Rain</title>
      <genre>Fantasy</genre>
      <publisher>John Doe</publisher>
    </book>
  </bookstore>
</catalog>
```

or

```
<?xml version="1.0"?>
<catalog>
  <bookstore>
    <shelf>
      <book id="bk102">
        <author>Ralls, Kim</author>
        <title>Midnight Rain</title>
        <genre>Fantasy</genre>
        <publisher>John Doe</publisher>
      </book>
    </shelf>
```

```

    </bookstore>
</catalog>

```

- Special characters such as hyphen (-) and period (.) in tag names are replaced by underscore (_) in the column name

XML tag names can contain hyphens and periods. Since the parser is converting XML to JSON, these characters become invalid object keys.

```

<?xml version="1.0"?>
<catalog>
  <book id="bk102">
    <author-name>Ralls, Kim</author-name>
    <title.1>Midnight Rain</title.1>
    <genre>Fantasy</genre>
    <publisher>John Doe</publisher>
  </book>
</catalog>

```

The generated columns are `id`, `author_name`, `title_1`, `genre`, and `publisher`.

- First-level only-text tags are ignored

```

<?xml version="1.0"?>
<catalog>
  <library> New Age Library </library>
  <book id="bk102">
    <author>Ralls, Kim</author>
    <title>Midnight Rain</title>
    <genre>Fantasy</genre>
    <publisher>John Doe</publisher>
  </book>
</catalog>

```

The `<library>` tag is ignored and only the content of the `<book>` tag is taken into account. The generated columns are `id`, `author`, `title`, `genre`, and `publisher`.

- First-level repetitive data is interpreted as an array of values

```

<?xml version="1.0" encoding="UTF-8"?>
<items id="orders">
  <item_number>1</item_number>
  <description>One Magic Christmas</description>
  <part>
    <unit_price>19.95</unit_price>
    <upccode>13131092899</upccode>
  </part>
  <quantity>9</quantity>
  <total>179.55</total>
  <item_number>2</item_number>
  <description>Lethal Weapon</description>
  <part>
    <unit_price>17.95</unit_price>
    <upccode>85391628927</upccode>
  </part>

```



```

    <quantity>5</quantity>
    <total>89.75</total>
</items>

```

The generated columns include `item_number`, `description`, `part`, and each column will have only one row with the following values respectively ([1,2], ["One Magic Christmas", "Lethal Weapon"], [{"unit_price":19.95,"upccode":13131092899}, {"unit_price":17.95,"upccode":85391628927}] and so on for the remaining columns.

- Tags containing values and attributes are converted to object

```

<?xml version="1.0"?>
<catalog>
  <book id="bk102">
    <author country="ca">Ralls, Kim</author>
    <title>Midnight Rain</title>
    <genre>Fantasy</genre>
    <publisher>John Doe</publisher>
  </book>
</catalog>

```

The `<author>` tag is converted to a column and will have an object as value that is structured as:

```

{
  "_": "Ralls, Kim",
  "country": "ca"
}


```

Note that the value of the tag has underscore ("`_`") as key and the attributes as "`attribute_name`": "`attribute_value`".

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 Database Actions page, click **Data Modeler**.
- Click **Selector**  to display the left navigation pane. Expand the **Development** menu and 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 3-11](#) shows a thumbnail preview that is displayed when you click the name of a diagram in the left pane.

Figure 3-11 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 SQL Developer Web 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 3-12](#) 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 3-12 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

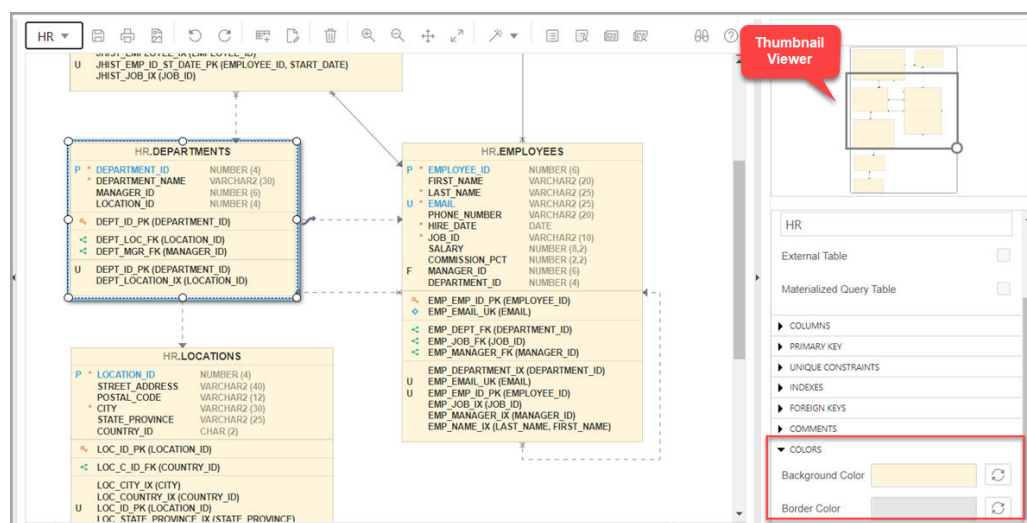
You can create and work with relational diagrams in the editor pane.

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 3-13](#) 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 3-13 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:

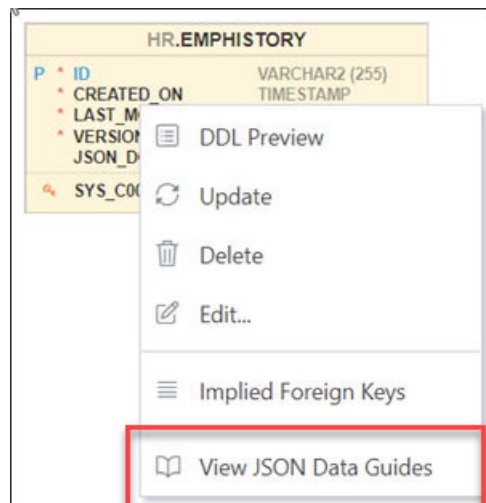
- **Save Diagram:** Saves the currently selected diagram. Diagrams are stored in a table that is created in the schema of the user.
- **Print Diagram:** Prints the selected diagram.
- **Save to SVG:** Saves the currently selected diagram to an image file in SVG format.
- **Add Objects to Diagram:** A 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" :
        {
```



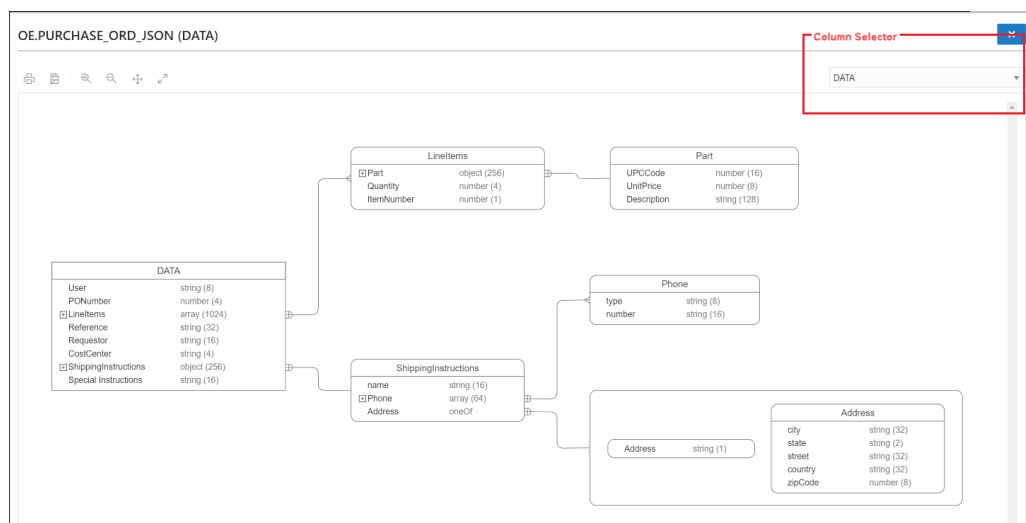
```
"type" :
{
  "type" : "string",
  "o:length" : 8,
  "o:preferred_column_name" : "DATA$type"
},
"number" :
{
  "type" : "string",
  "o:length" : 16,
  "o:preferred_column_name" : "DATA$number"
}
}
},
"Address" :
{
  "oneOf" :
  [
    {
      "type" : "string",
      "o:length" : 1,
      "o:preferred_column_name" : "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 3-14 JSON Data Guide Diagram



See Also:

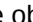
[JSON Data Guide in *Oracle Database JSON Developer's Guide*](#)

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](#)

If you do not have any existing objects for a object type (such as tables), you can create objects by using the **Object submenu**  icon, after first selecting the object type from the drop-down list. For all objects, the context menu includes **Quick DDL**, which saves the DDL statements for creating the object to a worksheet or to a text file. The various object types and the corresponding options are:

- **Tables:** Tables are used to hold data. Each table has multiple columns that describe the attributes of the database entity associated with the table, and each column has an associated data type.

You can create new tables, edit existing tables, and add indexes by using the options available in the context menu. Or you can select **Use as Template**, which displays the dialog box with information for the selected object filled in. You should then change the object name and any other aspects of the definition, as needed, to create a new object of that type.
- **Views:** Views are virtual tables that select data from one or more underlying tables.
- **Indexes:** An index contains an entry for each value that appears in the indexed column of the table and provides direct, fast access to rows.
- **Packages, Functions, Procedures:** A package contains subprograms, which are programming objects that can be stored and executed in the database server, and called from other programming objects or applications. A package can contain functions or procedures, or both. Select **Run** from the context menu to insert the code for a selected function, procedure, or package into the worksheet.
- **Triggers:** Triggers are stored PL/SQL blocks associated with a table, a schema, the database, or anonymous PL/SQL blocks or calls to a procedure implemented in PL/SQL or Java.
- **Types:** A data type associates a fixed set of properties with the values that can be used in a column of a table or in an argument of a function or procedure.
- **Sequences:** Sequences are used to generate unique integers. You can use sequences to automatically generate primary key values.
- **Materialized View:** A materialized view is a database object that contains the results of a query. The FROM clause of the query can name tables, views, and other materialized views.
- **Materialized Views Logs:** A materialized view log is a table associated with the master table of a materialized view. When DML changes are made to master table data, Oracle Database stores rows describing those changes in the materialized view log and then uses the materialized view log to refresh materialized views based on the master table.

- **Synonyms, Public Synonyms:** Synonyms provide alternative names for tables, views, sequences, procedures, stored functions, packages, materialized views, Java class database objects, user-defined object types, or other synonyms.
- **Database Links:** A database link is a database object in one database that enables you to access objects on another database. After you have created a database link, you can use it to refer to tables and views in the other database.

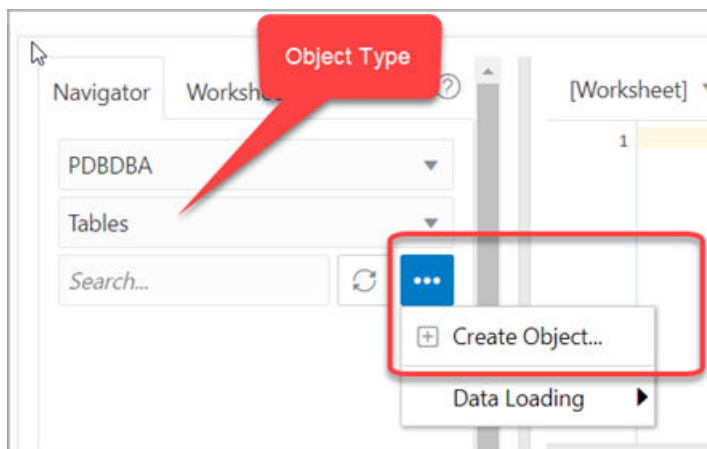
3.4.1 The Table Properties Dialog

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.

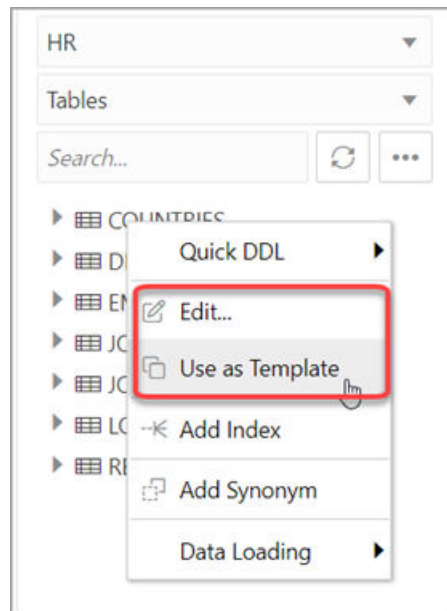
To create a table for a specific schema, in the Navigator tab, select **Tables** from the object type drop-down list, click **Object submenu** **...**, and select **Create Object**.

Figure 3-15 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 3-16](#).

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

Figure 3-16 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.

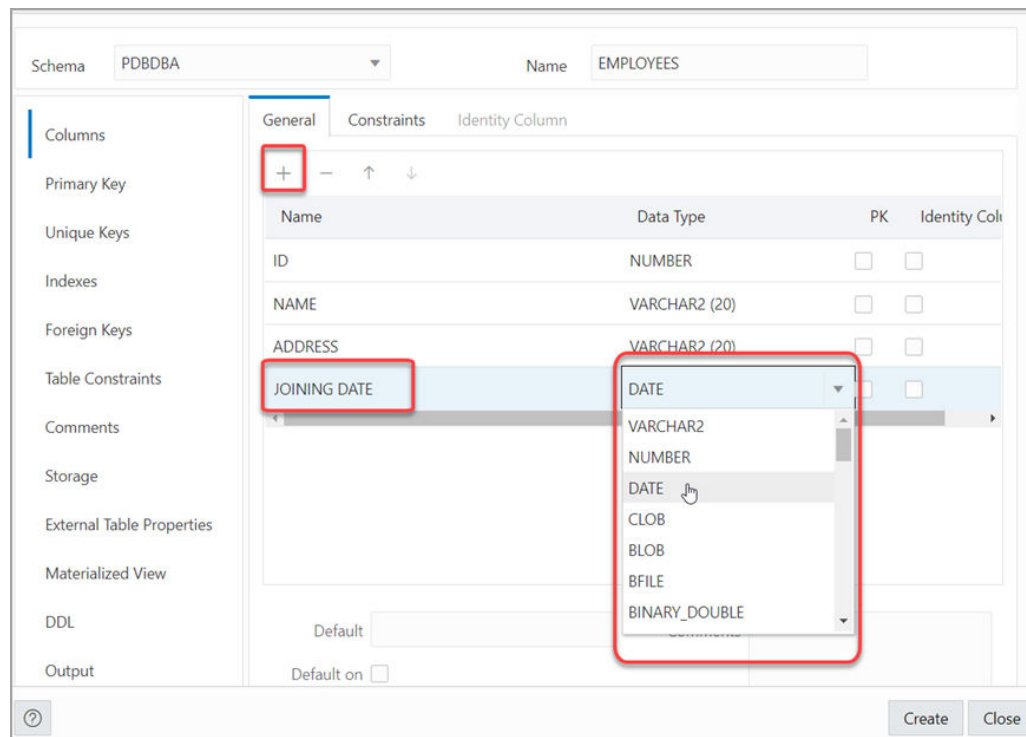
General tab

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 3-17 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 AMERICAN
  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 AMERICAN
    BADFILE log_file_dir:'ext_lv3.bad'
    LOGFILE log_file_dir:'ext_lv3.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":'salelv3.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 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:

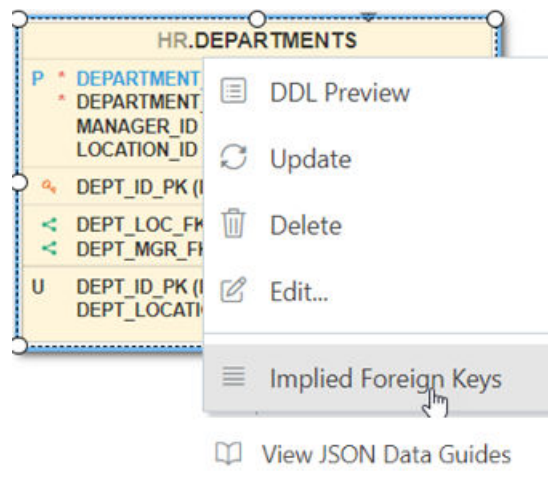
- Using the Implied Foreign Keys dialog
- By dragging the arrow to the referenced object in the diagram

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 3-18 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 3-19 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 3-20 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.

3.4.7 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.

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 Database Actions page, click **REST**.
- Click **Selector**  to display the left navigation pane. Expand the **Development** menu and select **REST**.

Topics

- [About RESTful Web Services](#)
- [RESTful Services Terminology](#)
- [About the Overview 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:

[Getting Started with RESTful Services in Oracle REST Data Services Installation, Configuration, and Development Guide](#)

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 Database Actions page, click **REST**.
- Click **Selector**  to display the left navigation pane. Expand the **Development** menu and select **REST**.

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

The Overview page consists of three sections.

Objects



Displays the total number of resource modules, roles, privileges, and OAuth clients created. Click the respective card to navigate to the corresponding pages for modules, roles, privileges and OAuth clients.

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


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 Oracle SQL Developer Web](#)

4.4 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.

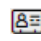
Filter by: Performs a search based on the filters selected. In the Filter by drop down, first select the column to apply the filter on, and then select the values to search for in that column.


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.

4.5 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](#)
- [Viewing Resource Handler Details and Managing Parameters](#)



See Also:

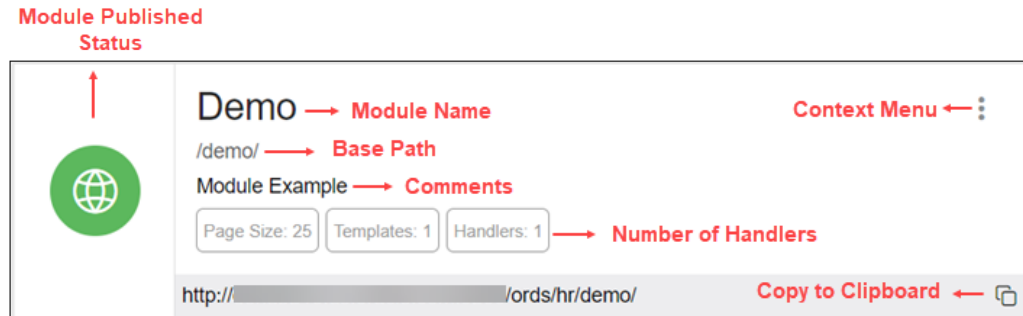
[Creating a RESTful Service in Oracle REST Data Services Installation, Configuration, and Development Guide](#)

4.5.1 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.



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](#)
- [Editing a Resource Module](#)
- [Deleting a Resource Module](#)
- [Publishing/Unpublishing a Resource Module](#)
- [Exporting a Resource Module](#)
- [Navigating to the Templates page for a module](#)

See Also:

[About the REST Search Toolbar](#)

4.5.1.1 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.
- **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`

3. Click **Create**.

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


4.5.1.2 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](#).


4.5.1.3 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.


4.5.1.4 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**.


4.5.1.5 Exporting a Resource Module

You can export the PL/SQL source code for a resource module or RESTful service using the Copy to Clipboard option. If the module has templates, handlers, parameters, roles, and so on, these objects are also displayed and available for copying to the clipboard.

1. In the Modules page, for a specific module, click  and select **Export**.
The PL/SQL code for the module appears.
2. Click **Copy to Clipboard**.


4.5.2 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

4.5.2.1 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.


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.

4.5.2.2 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](#).


4.5.2.3 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.

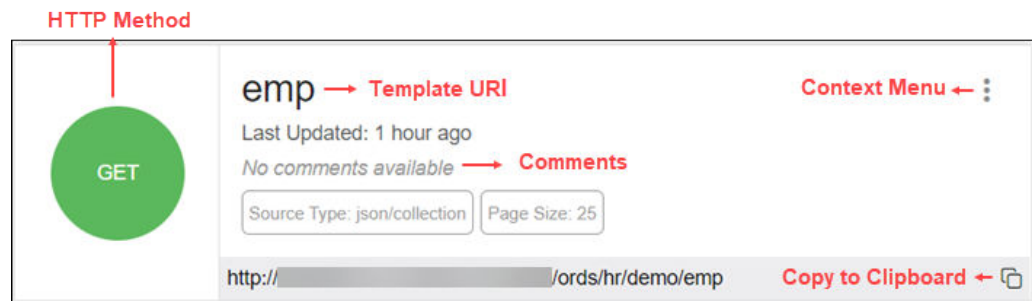
4.5.3 Managing Resource Handlers

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.

Click the handler name to go to the HTTP method page or you can also navigate to the HTTP method page using the context menu.

The actions available in the context menu are:

- [Creating a Resource Handler](#)

 **Note:**

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

- [Editing a Resource Handler](#)
- [Deleting a Resource Handler](#)
- Navigating to the Handler *HTTP method* page

4.5.3.1 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.

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.



See Also:

[About cURL and Testing RESTful Services in Oracle REST Data Services Installation, Configuration, and Development Guide](#)

4.5.3.2 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:

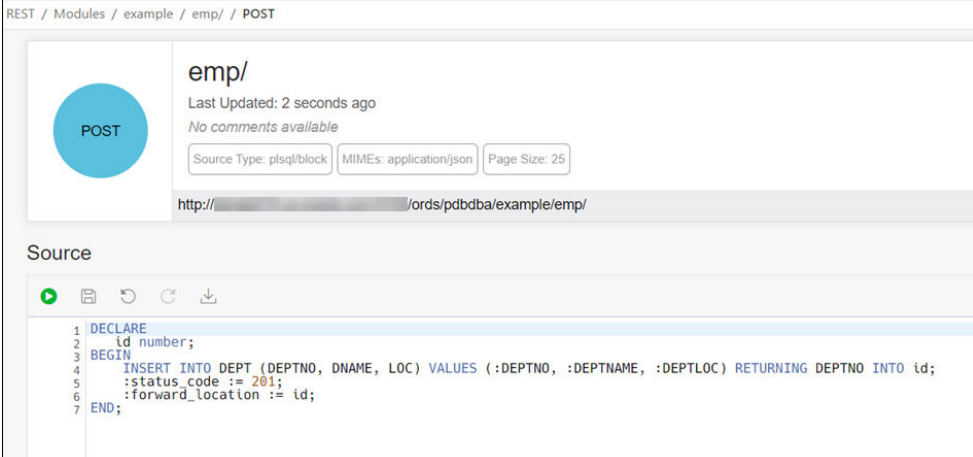
1. In the Handlers page, click **Create Handler**.
2. Enter the following details:
 - In the **Method** field, select **POST**.
 - In the **Source Type** field, select **PL/SQL**.
 - In the **Source** field, enter the following PL/SQL block:

```
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.



REST / Modules / example / emp / POST

emp/
Last Updated: 2 seconds ago
No comments available

Source Type: plsql/block | MIMEs: application/json | Page Size: 25

http://.../ords/pdbdba/example/emp/

Source

```
1 DECLARE
2   id number;
3 BEGIN
4   INSERT INTO DEPT (DEPTNO, DNAME, LOC) VALUES (:DEPTNO, :DEPTNAME, :DEPTLOC) RETURNING DEPTNO INTO id;
5   :status_code := 201;
6   :forward_location := id;
7 END;
```

You can test the source code by clicking **Execute**  in the Source editor.

3. 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
```



4. 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;
```

4.5.3.3 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](#).


4.5.3.4 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.

4.5.4 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.
- 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](#)

See Also:

[About the REST Search Toolbar](#)

4.5.4.1 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.
3. Click **Create**.

The parameter is added in the Parameters section.

4.5.4.2 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](#).

4.5.4.3 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 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](#)

**See Also:**

[Tutorial for Protecting and Accessing Resources in Oracle REST Data Services](#)

4.6.1 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 for Accessing Resources in Oracle REST Data Services Installation, Configuration, and Development Guide](#)

4.6.1.1 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.
3. Click **Create**.

The new assigned role is displayed on the Roles page.

4.6.1.2 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](#).


4.6.1.3 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.

4.6.1.4 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.

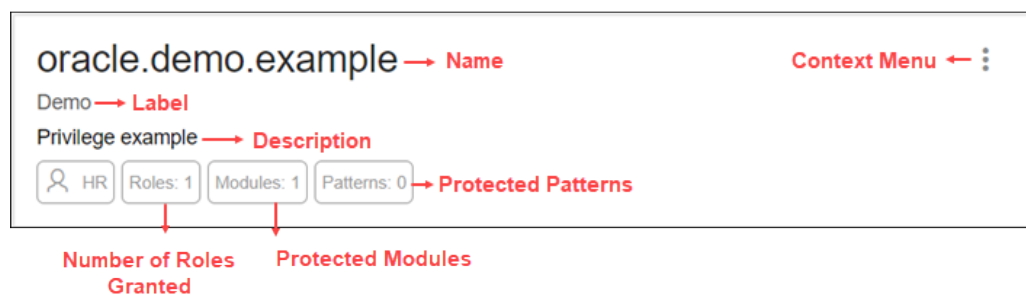
4.6.2 Managing Privileges

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 actions available in the context menu are:

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

4.6.2.1 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.
3. Click **Create**.

The new privilege is displayed on the Privileges page.

4.6.2.2 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](#).

4.6.2.3 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**.

4.6.3 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.



The actions available in the context menu are:

- [Creating an OAuth Client](#)
- [Editing an OAuth Client](#)
- [Deleting an OAuth Client](#)
- **Get Bearer Token:** Provides the access token to call the RESTful service.



See Also:


[Tutorial for OAuth Flow in Oracle REST Data Services](#)

4.6.3.1 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:
 - **Name:** Name of the client.
 - **Description:** Description of the purpose of the client.
 - **Support URI:** Enter the URI where end users can contact the client for support. Example: http:// www.myclientdomain.com/support/
 - **Support Email:** Enter the email where end users can contact the client for support.
 - **Roles:** Select roles to be granted.
 - **Allowed Origins:** Add the list of URL prefixes.
 - **Privileges:** Select privileges that the client wants to access.
3. Click **Create**.

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

The Client ID and Client Secret values represent the secret credentials 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.

 **See Also:**

[Tutorial for Protecting and Accessing Resources in Oracle REST Data Services Installation, Configuration, and Development Guide](#)

4.6.3.2 Example: Creating an OAuth Client

This section describes how to create an OAuth Client using an example.

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:
 - Enter Name, Description, 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 AMccw1nt99gm9kxDs_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;
```

4.6.3.3 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](#).

4.6.3.4 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**.

5


The JSON Page



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 SQL Developer Web 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 Database Actions page, click **JSON**.
- Click **Selector**  to display the left navigation pane. Expand the **Development** menu and select **JSON**.

Topics

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



See Also:

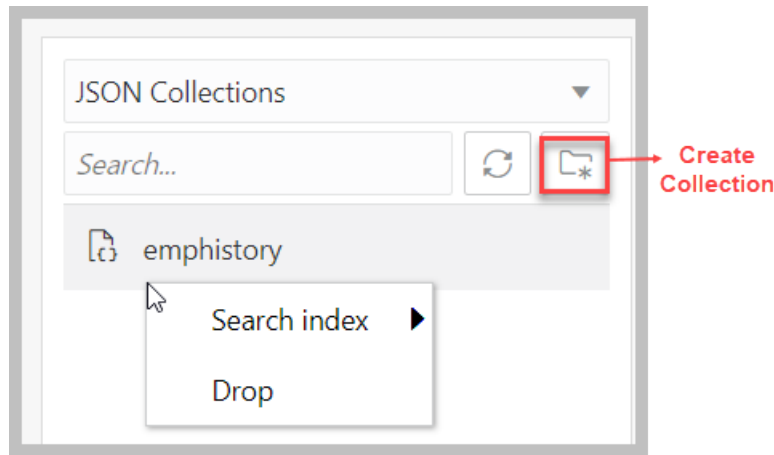
[SODA Overview in Oracle Database Introduction to Simple Oracle Document Access \(SODA\)](#)

5.1 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.

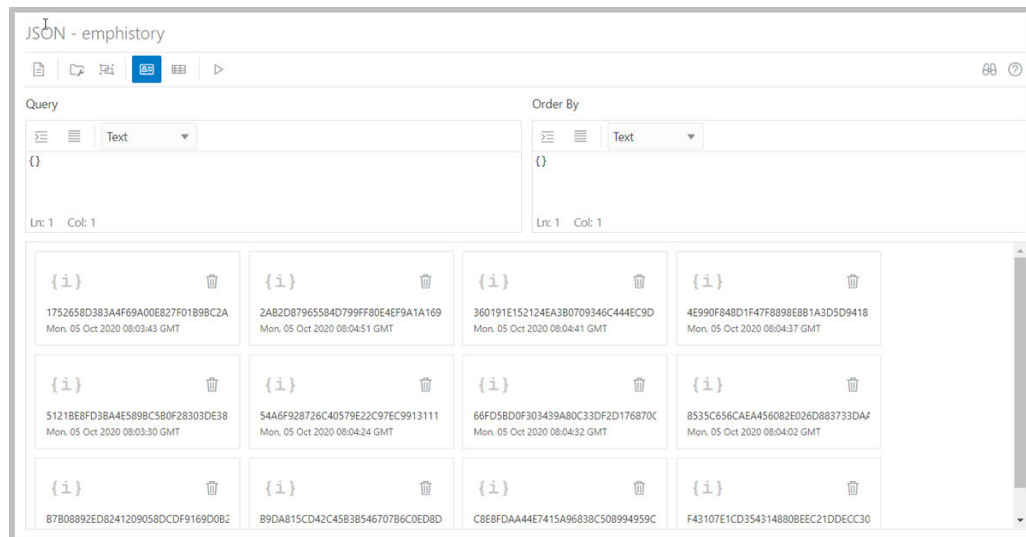
Click the Create Collection icon to create a new collection. See [Creating a Collection](#)

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




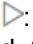
- **Search index:** Create a JSON search index for the collection. You also have the option to drop the created search index. See [Managing JSON Search Indexes for a Collection](#)
- **Drop:** Removes the collection from the database completely.

Viewing the Contents of a Collection

When a specific collection is selected in the left pane, the contents of the collection are displayed in the right pane.




The icons in the toolbar are:

- **New JSON Document** : Adds a new document to the collection. See [About Adding or Editing a JSON Document](#)
- **Collection Details** : Displays the properties of a collection. For a description of the fields, see [Creating a Collection](#).
- **Diagram** : Displays the the JSON data guide as a diagram in a hierarchical format. See [Viewing the JSON Data Guide Diagram for a Collection](#)
- There are two views available for documents. In both views, each document entry displays a unique document identifier and the created timestamp. Click the delete icon to delete a document.
 - **Card View**: Displays each document in a card format. This is the default view.
 - **List View**: Displays the documents in a list.
- **Filter** : Filters documents using the QBE condition entered in the Query or Order By field. See [About Querying Documents in a Collection](#)


5.2 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.
- **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.

5.2.1 Creating a Collection

This section describes how to create a collection.

1. In the left pane, click **Create Collection** .
2. Enter the collection name and click **Create**. A brief description of each field is mentioned below:

Collection Name: Enter a name for the collection. Note that the collection name is case sensitive. By default, this is the only field where you need to make an entry to create a collection. All others fields have default pre-filled values.

General

- **Search Index**: Enter a name to create the JSON search index for the collection.

- **Schema:** Name of the Oracle Database schema that owns the table or view to which the collection is mapped.
- **Table:** Name of the table or view to which the collection is mapped.

Key Column

- **Column Name:** Name of the column that stores the document key. The default value is ID.
- **Data Type:** SQL data type of the column that stores the document key. The default value is VARCHAR2.
- **Max Length:** Maximum length of the column in bytes (for type VARCHAR2 only). The default value is 255.
- **Assignment Method:** Method used to assign keys to objects that are inserted into the collection. The default value is UUID. If you specify the method as SEQUENCE, then you must also specify the name of that sequence.

Content Column


- **Column Name:** Name of the column that stores the document content. The default value is JSON_DOCUMENT.
 - **Data Type:** SQL data type of the column that stores the document content. The default value is BLOB.
 - **Max Length:** Maximum length of the content column in bytes (for type VARCHAR2 only). The default value is 4000.
 - **JSON Validation:** Syntax to which JSON content must conform: Standard, Strict or Lax.
 - **LOB Compression:** Specifies the SecureFiles LOB compression setting.
 - **LOB Cache:** Specifies the SecureFiles LOB cache setting. The default is enabled.
 - **LOB Encryption:** Specifies the SecureFiles LOB encryption setting.
 - **Version Column:** Name of the column that stores the document version. The default value is VERSION.
 - **Generation Method:** Method used to compute version values for objects when they are inserted into a collection or replaced. The default value is SHA256.
 - **Media Type Column:** Not supported in the 20.3 release.
 - **Creation Time Column:** Name of the column that stores the creation time stamp of the document. The default value is CREATED_ON.
 - **Last Modified Column:** Name of the column that stores the last-modified time stamp of the document. The default value is LAST_MODIFIED.
 - **Index:** Name of the index on the last-modified column. The index is created if a name is specified.
3. Click **Create** to create the collection.
Click **Cancel** or press the **Esc** key to cancel the collection.

See Also:

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

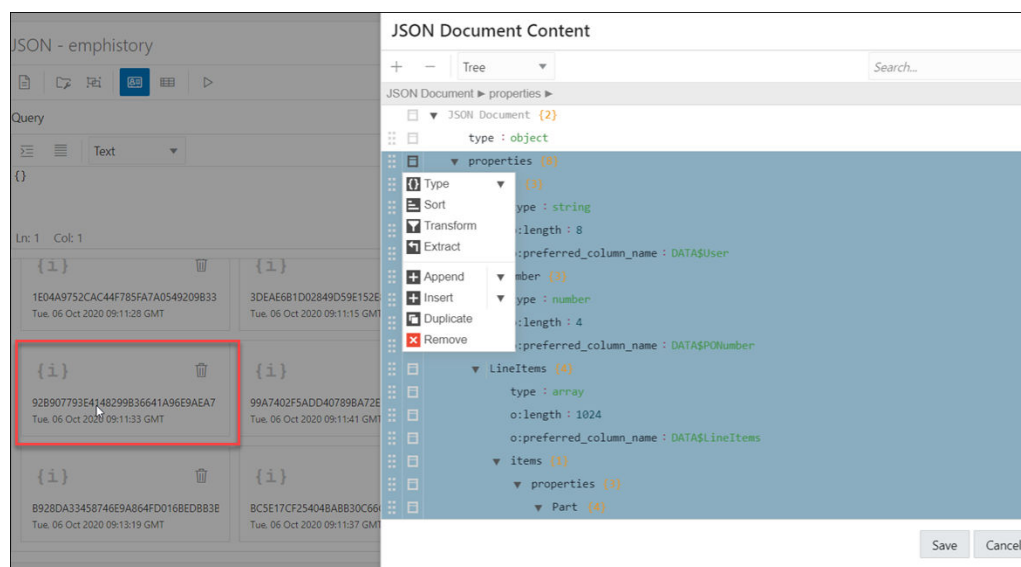
5.2.2 About Adding or Editing a JSON Document

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.

To edit a document, click the specific document in the collection.

The following figure shows a document edited in the Tree mode in the JSON editor.



The JSON editor enables you to view, edit, format and validate JSON documents. The various modes in the editor are:

- **Form:** Modify values but not the data structure.
- **View:** Cannot modify data structure or values.
- **Text:** Use to copy or type content for a new JSON document.
- **Tree:** Add, change, remove, move and duplicate fields and values.

For form, view and tree modes, you can use the Search field in the top right corner to search for and highlight matching text in the document.

5.3 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)*.

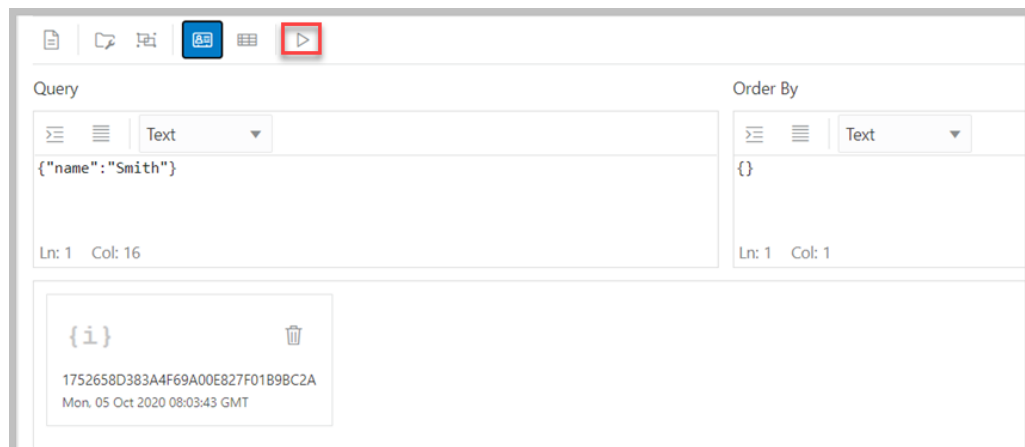
For a specific collection, you can enter the QBE filter in the Query or Order By fields in the right pane. The Query field is used to enter the filter condition and the Order By field is used to sort query results in ascending or descending order. You can use these fields separately or together to filter the documents.

Both fields are JSON editors, where you can select the mode and enter the condition that you want to use for filtering the documents.

There are two modes available in the editor: Text and Tree.

Text Mode

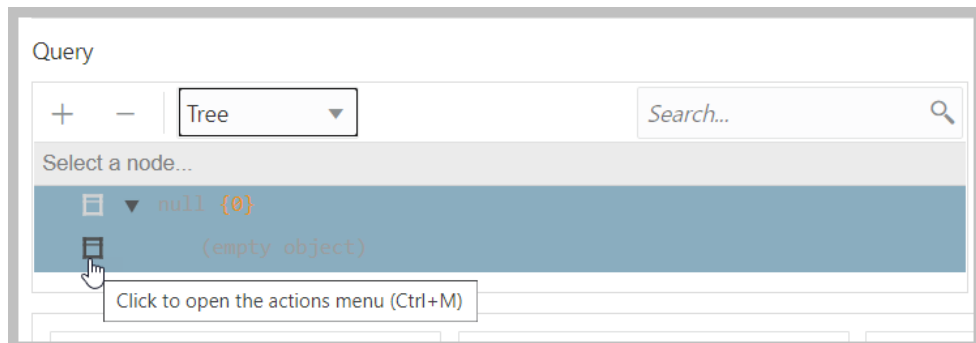
Enter the filter condition in the Query field. For example, to select documents where the name is Smith, enter `{"name": "Smith"}` and then click **Filter** ▶.



Tree Mode

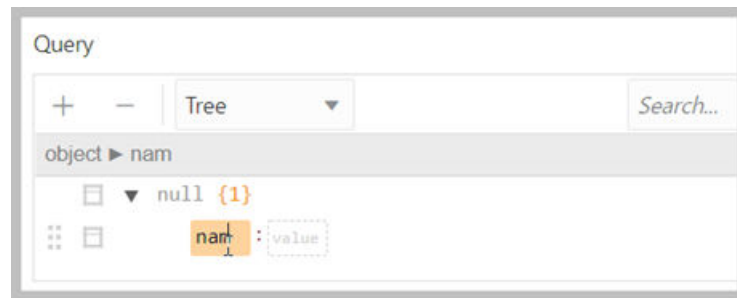
In the Tree mode:

1. Click the box to the left of (empty object).



2. Click the + symbol before Append.

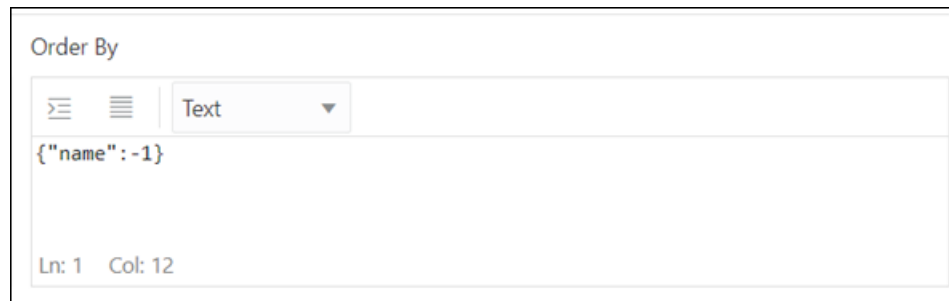
3. Enter the field name and value.



4. Click **Filter** ▶

Sort Results in Text Mode Using Order By

To sort the names of employees in descending order, use the expression as shown in the following figure. In the Order By field, a value of 1 specifies ascending order while a minus symbol (-1) specifies descending order.




5.4 Managing JSON Search Indexes for a Collection

You can create or drop a JSON search index for a collection.

Create a JSON Search Index

1. In the left pane, right-click the specific collection, select **Search Index** and then select **Create**.
2. Enter a name for the search index and click **OK**.

The search index name is displayed when viewing collection details.



Drop a Search Index

1. In the left pane, right-click the specific collection, select **Search Index** and then select **Drop**.

You are prompted to confirm.


2. Click **OK** to drop.

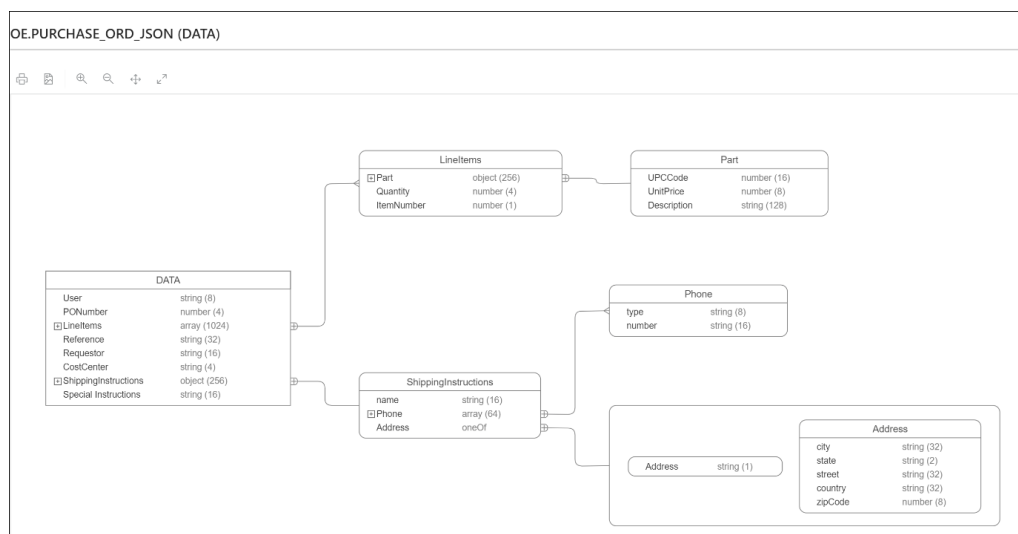
5.5 Viewing the JSON Data Guide Diagram for a Collection

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*

6


The Database Users Page



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 users.

To navigate to the Database Users page, do either of the following:


- In the Database Actions page, click **Database Users**.
- Click **Selector**  to display the left navigation pane. Expand the **Administration** menu and 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 SQL Developer Web user login page. Use **Copy to Clipboard**  to copy the URL to the user's clipboard.

Click **Object submenu**  to open the context menu. The actions available are:

- **Edit**: Opens the Edit User Dialog, where you can edit current user information. See [About Create/Edit 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 [About REST Enable User](#).
- **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 [About Delete 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 [About Create/Edit 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.

6.1 About Create/Edit 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.
- **New 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.
- **REST Enable:** If this option is selected, the user is enabled for REST. Enter the related fields: Alias, URL Mapping Type and Authorization required. These options are also available in the REST Enable dialog.

Granted Roles Tab

Specifies roles to be granted to the user. For each role, you can check **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 REST Enable is selected.

6.2 About REST Enable User

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.

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.

6.3 About Delete User

Use this option to delete users.

In the Database Users page, select **Delete** from the user's context menu to delete a user.

- **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 or error notification is displayed.

7

The Monitoring Pages



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

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](#)

7.1 The Monitoring Overview Page



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 pane, click the down arrow to expand **Monitoring**, and then click **Overview**.

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.


7.2 The Performance Hub Page



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 Database Actions page, click **Performance Hub**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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.

Time Range Area

Use these controls in the time range area to specify the time period for which you want to view performance data:

- **Select Duration:** Use this drop-down list, located in the top right of the time range area, to set the timeframe displayed in the timeline. You can choose **Last hour**, **Last 8 hours**, **Last 24 hours**, **Last Week**, or you can choose **Custom** and define your own timeframe.
- **Timeline:** The timeline displays a graph spanning the timeframe selected in the timeframe dropdown, showing Waits, User I/O and CPU usage during the period. At its end is the time slider.
- **Time Slider:** The time slider is a box you can drag back and forth along the current timeline. Use it to pick the specific time period within the timeframe for which you want to view performance data. You can also drag the side handles on the time slider to make it wider or narrower to encompass a longer or shorter time period.

ASH Analytics Tab

The ASH Analytics tab consists of the Average Active Sessions chart and two secondary tables below it.

- **Average Active Sessions Chart:** This chart shows performance information for the time period defined by the time slider. You can choose to chart different dimensions of information, such as **Wait Class**, **Wait Event**, or **Service**, by selecting the dimension from the drop-down list to the right of the chart title.

You can download an AWR (Automatic Workload Repository) report for the current time period by right-clicking in the Average Active Sessions chart area and choosing **Generate AWR Report**.

For more information about ASH and AWR, see Active Session History (ASH) and Automatic Workload Repository (AWR) in *Oracle Database Concepts*.

- **Secondary Tables:** The two tables below the Average Active Sessions chart show the information dimension chosen in the chart filtered by another dimension you choose. For example, if the Average Active Sessions chart is showing Wait Class, you could show SQL ID and User Session dimension information by Wait Class, one in each of the two secondary tables.

In the secondary tables, the SQL ID and User Session dimensions provide links to SQL Details and Session Details pages, respectively, for the dimension data listed in the table.

SQL Monitoring Tab

The SQL Monitoring tab shows a table of the top 100 monitored SQL statements that were executing or that completed during the selected time period.

The table displays information about monitored SQL statement executions. If there is a green spinning icon in the Status column, then the monitored statement did not complete during the selected time period. A red cross indicates that the SQL did not complete either due to an error or due to the session getting terminated. If there is a

check mark in the Status column, then the statement completed its execution during the selected time period.

SQL statements are monitored only if they have consumed at least 5 seconds of CPU or I/O time.

You can view information such as the status of a statement, its duration, its type (SQL, PL/SQL, or DBOP), its SQL ID, its SQL plan hash, the user who issued it, whether it executed as a serial or parallel statement, the time the database spent performing CPU activity, I/O, or other activity for the statement, the read and write requests and bytes associated with the statement, and the start and end time for the statement.

Click a SQL ID to display the SQL Details page with more information about that SQL statement.


7.3 The Instance Viewer Page



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 Database Actions page, click **Instance Viewer**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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 7-1 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)

7.4 The Logins Page




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 Database Actions page, click **Logins**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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 7-2 Tabs in Last Hour Statistics

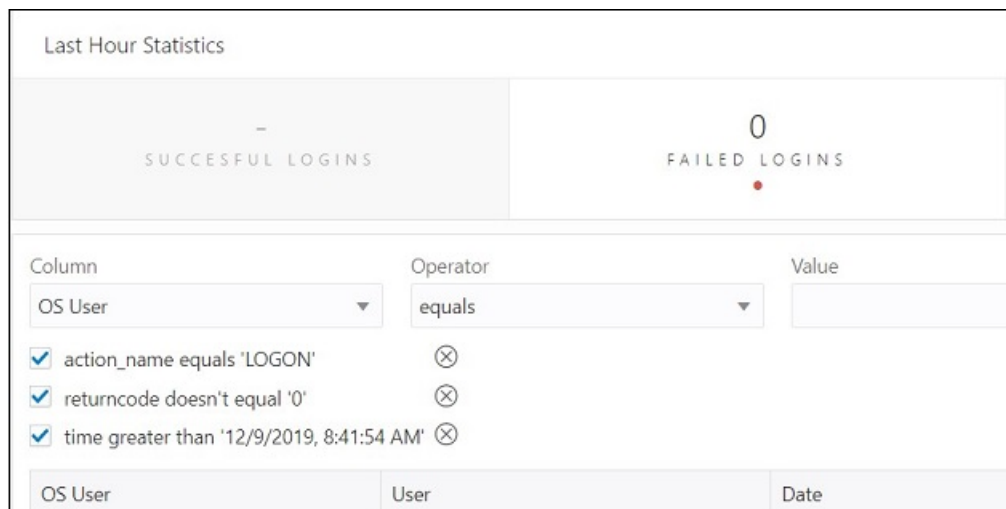





You can search for results in the table by entering values in the Column, Operator and Value fields and clicking Filter . You can also filter by selecting one or more of the available filter options for each action.

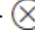
Click Refresh  to refresh the data in the grid.

The following figure shows the three filters for the Failed Logins action.

Figure 7-3 Filters for Failed Logins



Column	Operator	Value
OS User	equals	
<input checked="" type="checkbox"/> action_name equals 'LOGON'		
<input checked="" type="checkbox"/> returncode doesn't equal '0'		
<input checked="" type="checkbox"/> 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.


7.5 The Alerts Page



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 Database Actions page, click **Alerts**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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**.


7.6 The Sessions Page



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 Database Actions page, click **Sessions**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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**.


7.7 The Storage Page



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 Database Actions page, click **Storage**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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.


7.8 The Parameters Page




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

The Parameters page 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 Database Actions page, click **Parameters**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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**.

7.9 The Real Time SQL Monitoring Page




Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles for Oracle Database 18c and previous 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 Database Actions page, click **Real Time SQL Monitor**.

- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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:

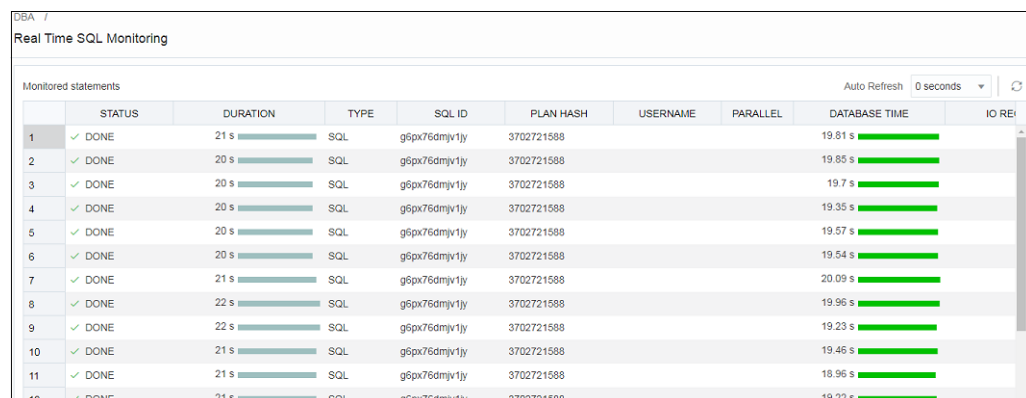
- General view of monitored statements
- View of SQL execution details

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 7-4 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	

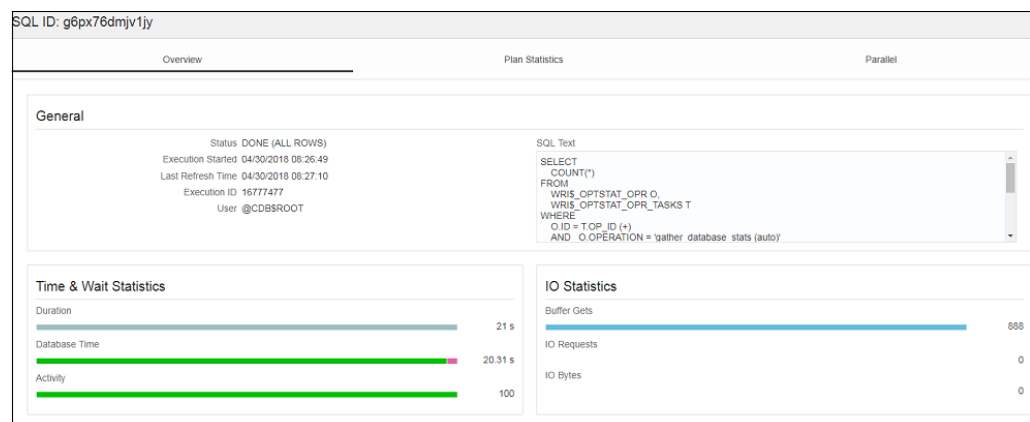
View of SQL Execution Details

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 7-5 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*.


7.10 The Top SQL Page



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 Database Actions page, click **Top SQL**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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.


7.11 The Waits Page



Available only if you signed in as a database user with **DBA** and **PDB_DBA** roles for Oracle Database 18c and previous releases.

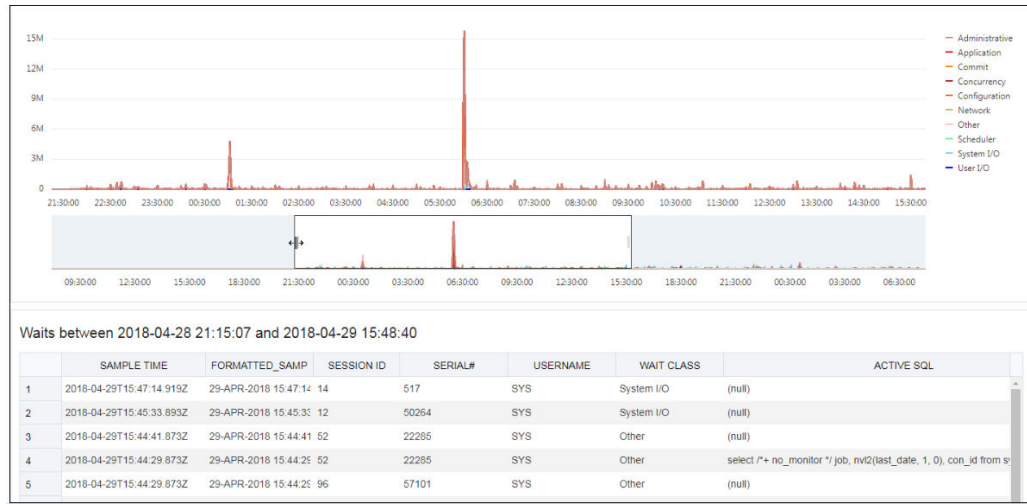
The Waits page shows a chart with the distribution of wait events in real time.

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

- In the Database Actions page, click **Waits**.
- Click **Selector**  to display the left navigation pane. Expand the **Monitoring** menu and 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 7-6 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 7-7 Display Wait Event Details



A

Supported SQL*Plus and SQLcl Commands

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] [printspect [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

Starting from SQL Developer Web release 19.4.2, SODA (Simple Oracle Document Access) commands are supported. 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>*