

Oracle® Database

JSON Developer's Guide



19c
E96227-08
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®

Oracle Database JSON Developer's Guide, 19c

E96227-08

Copyright © 2015, 2020, Oracle and/or its affiliates.

Primary Author: Drew Adams

Contributors: Oracle JSON development, product management, and quality assurance teams.

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	xiv
Documentation Accessibility	xiv
Related Documents	xiv
Conventions	xv
Code Examples	xvi

Changes in This Release for Oracle Database JSON Developer's Guide

Changes in Oracle Database Release 19c, Version 19.1, for JSON Developer's Guide	xvii
Changes in Oracle Database Release 18c, Version 18.1, for JSON Developer's Guide	xx

Part I Introduction to JSON Data and Oracle Database

1 JSON in Oracle Database

1.1 Overview of JSON in Oracle Database	1-2
1.2 Getting Started Using JSON with Oracle Database	1-4
1.3 Oracle Database Support for JSON	1-6

2 JSON Data

2.1 Overview of JSON	2-1
2.2 JSON Syntax and the Data It Represents	2-2
2.3 JSON Compared with XML	2-4

Part II Store and Manage JSON Data

3	Overview of Storing and Managing JSON Data	
4	Creating a Table With a JSON Column	
4.1	Determining Whether a Column Necessarily Contains JSON Data	4-3
5	SQL/JSON Conditions IS JSON and IS NOT JSON	
5.1	Unique Versus Duplicate Fields in JSON Objects	5-2
5.2	About Strict and Lax JSON Syntax	5-2
5.3	Specifying Strict or Lax JSON Syntax	5-5
6	Character Sets and Character Encoding for JSON Data	
7	Partitioning JSON Data	
8	Replication of JSON Data	
Part III Insert, Update, and Load JSON Data		
9	Overview of Inserting, Updating, and Loading JSON Data	
10	Loading External JSON Data	
11	Updating a JSON Document with JSON Merge Patch	
Part IV Query JSON Data		
12	Simple Dot-Notation Access to JSON Data	

13 SQL/JSON Path Expressions

13.1	Overview of SQL/JSON Path Expressions	13-1
13.2	SQL/JSON Path Expression Syntax	13-2
13.2.1	Basic SQL/JSON Path Expression Syntax	13-2
13.2.2	SQL/JSON Path Expression Syntax Relaxation	13-9
13.3	SQL/JSON Path Expression Item Methods	13-11
13.4	ISO 8601 Date and Time Support	13-14
13.5	Types in Comparisons	13-16

14 Clauses Used in SQL Query Functions and Conditions

14.1	RETURNING Clause for SQL Query Functions	14-1
14.2	Wrapper Clause for SQL/JSON Query Functions JSON_QUERY and JSON_TABLE	14-4
14.3	Error Clause for SQL Query Functions and Conditions	14-5
14.4	Empty-Field Clause for SQL/JSON Query Functions	14-8
14.5	ON MISMATCH Clause for JSON_VALUE	14-9

15 SQL/JSON Condition JSON_EXISTS

15.1	Using Filters with JSON_EXISTS	15-2
15.2	JSON_EXISTS as JSON_TABLE	15-4

16 SQL/JSON Function JSON_VALUE

16.1	Using SQL/JSON Function JSON_VALUE With a Boolean JSON Value	16-3
16.2	SQL/JSON Function JSON_VALUE Applied to a null JSON Value	16-4
16.3	Using JSON_VALUE To Instantiate a User-Defined Object Type Instance	16-4
16.4	JSON_VALUE as JSON_TABLE	16-7

17 SQL/JSON Function JSON_QUERY

17.1	JSON_QUERY as JSON_TABLE	17-2
------	--------------------------	------

18 SQL/JSON Function JSON_TABLE

18.1	SQL NESTED Clause Instead of JSON_TABLE	18-4
18.2	COLUMNS Clause of SQL/JSON Function JSON_TABLE	18-5
18.3	JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions	18-8
18.4	Using JSON_TABLE with JSON Arrays	18-9

18.5	Creating a View Over JSON Data Using JSON_TABLE	18-12
------	---	-------

19 Oracle SQL Function JSON_SERIALIZE

20 JSON Data Guide

20.1	Overview of JSON Data Guide	20-2
20.2	Persistent Data-Guide Information: Part of a JSON Search Index	20-4
20.3	Data-Guide Formats and Ways of Creating a Data Guide	20-7
20.4	JSON Data-Guide Fields	20-9
20.5	Data-Dictionary Views For Persistent Data-Guide Information	20-12
20.6	Specifying a Preferred Name for a Field Column	20-13
20.7	Creating a View Over JSON Data Based on Data-Guide Information	20-15
20.7.1	Creating a View Over JSON Data Based on a Hierarchical Data Guide	20-16
20.7.2	Creating a View Over JSON Data Based on a Path Expression	20-19
20.8	Adding and Dropping Virtual Columns For JSON Fields Based on Data-Guide Information	20-23
20.8.1	Adding Virtual Columns For JSON Fields Based on a Hierarchical Data Guide	20-24
20.8.2	Adding Virtual Columns For JSON Fields Based on a Data Guide-Enabled Search Index	20-27
20.8.3	Dropping Virtual Columns for JSON Fields Based on Data-Guide Information	20-30
20.9	Change Triggers For Data Guide-Enabled Search Index	20-30
20.9.1	User-Defined Data-Guide Change Triggers	20-32
20.10	Multiple Data Guides Per Document Set	20-35
20.11	Querying a Data Guide	20-38
20.12	A Flat Data Guide For Purchase-Order Documents	20-41
20.13	A Hierarchical Data Guide For Purchase-Order Documents	20-47

Part V Generation of JSON Data

21 Generation of JSON Data with SQL/JSON Functions

21.1	Overview of SQL/JSON Generation Functions	21-1
21.2	Handling of Input Values For SQL/JSON Generation Functions	21-3
21.3	SQL/JSON Function JSON_OBJECT	21-6
21.4	SQL/JSON Function JSON_ARRAY	21-10
21.5	SQL/JSON Function JSON_OBJECTAGG	21-12

Part VI PL/SQL Object Types for JSON

22 Overview of PL/SQL Object Types for JSON

23 Using PL/SQL Object Types for JSON

Part VII GeoJSON Geographic Data

24 Using GeoJSON Geographic Data

Part VIII Performance Tuning for JSON

25 Overview of Performance Tuning for JSON

26 Indexes for JSON Data

26.1	Overview of Indexing JSON Data	26-2
26.2	How To Tell Whether a Function-Based Index for JSON Data Is Picked Up	26-3
26.3	Creating Bitmap Indexes for SQL/JSON Condition JSON_EXISTS	26-3
26.4	Creating JSON_VALUE Function-Based Indexes	26-3
26.5	Using a JSON_VALUE Function-Based Index with JSON_TABLE Queries	26-5
26.6	Using a JSON_VALUE Function-Based Index with JSON_EXISTS Queries	26-6
26.7	Data Type Considerations for JSON_VALUE Indexing and Querying	26-8
26.8	Indexing Multiple JSON Fields Using a Composite B-Tree Index	26-9
26.9	JSON Search Index: Ad Hoc Queries and Full-Text Search	26-10

27 In-Memory JSON Data

27.1	Overview of In-Memory JSON Data	27-1
27.2	Populating JSON Data Into the In-Memory Column Store	27-3

A Oracle Database JSON Restrictions

B Diagrams for Basic SQL/JSON Path Expression Syntax

Index

List of Examples

2-1	A JSON Object (Representation of a JavaScript Object Literal)	2-3
4-1	Using IS JSON in a Check Constraint to Ensure JSON Data is Well-Formed	4-1
4-2	Inserting JSON Data Into a VARCHAR2 JSON Column	4-2
5-1	Using IS JSON in a Check Constraint to Ensure JSON Data is Strictly Well-Formed (Standard)	5-5
7-1	Creating a Partitioned Table Using a JSON Virtual Column	7-1
9-1	Inserting JSON Data Into a BLOB Column	9-1
10-1	Creating a Database Directory Object for Purchase Orders	10-2
10-2	Creating an External Table and Filling It From a JSON Dump File	10-2
10-3	Creating a Table With a BLOB JSON Column	10-2
10-4	Copying JSON Data From an External Table To a Database Table	10-2
11-1	A JSON Merge Patch Document	11-3
11-2	A Merge-Patched JSON Document	11-3
11-3	Updating a JSON Column Using JSON Merge Patch	11-3
11-4	Updating Selected JSON Documents On the Fly	11-4
12-1	JSON Dot-Notation Query Compared With JSON_VALUE	12-3
12-2	JSON Dot-Notation Query Compared With JSON_QUERY	12-3
15-1	JSON_EXISTS: Path Expression Without Filter	15-2
15-2	JSON_EXISTS: Current Item and Scope in Path Expression Filters	15-2
15-3	JSON_EXISTS: Filter Conditions Depend On the Current Item	15-3
15-4	JSON_EXISTS: Filter Downscoping	15-3
15-5	JSON_EXISTS: Path Expression Using Path-Expression exists Condition	15-3
15-6	JSON_EXISTS Expressed Using JSON_TABLE	15-4
16-1	JSON_VALUE: Returning a JSON Boolean Value to PL/SQL as BOOLEAN	16-3
16-2	JSON_VALUE: Returning a JSON Boolean Value to SQL as VARCHAR2	16-3
16-3	Instantiate a User-Defined Object Instance From JSON Data with JSON_VALUE	16-5
16-4	Instantiate a Collection Type Instance From JSON Data with JSON_VALUE	16-6
16-5	JSON_VALUE Expressed Using JSON_TABLE	16-7
17-1	Selecting JSON Values Using JSON_QUERY	17-2
17-2	JSON_QUERY Expressed Using JSON_TABLE	17-3
18-1	Equivalent JSON_TABLE Queries: Simple and Full Syntax	18-2
18-2	Equivalent: SQL NESTED and JSON_TABLE with LEFT OUTER JOIN	18-4
18-3	Accessing JSON Data Multiple Times to Extract Data	18-9
18-4	Using JSON_TABLE to Extract Data Without Multiple Parses	18-9
18-5	Projecting an Entire JSON Array as JSON Data	18-10

18-6	Projecting Elements of a JSON Array	18-10
18-7	Projecting Elements of a JSON Array Plus Other Data	18-10
18-8	JSON_TABLE: Projecting Array Elements Using NESTED	18-11
18-9	Creating a View Over JSON Data	18-13
18-10	Creating a Materialized View Over JSON Data	18-13
19-1	Using JSON_SERIALIZE To Convert BLOB Data To Pretty-Printed Text	19-1
20-1	Enabling Persistent Support for a JSON Data Guide But Not For Search	20-6
20-2	Disabling JSON Data-Guide Support For an Existing JSON Search Index	20-7
20-3	Gathering Statistics on JSON Data Using a JSON Search Index	20-7
20-4	Specifying Preferred Column Names For Some JSON Fields	20-14
20-5	Creating a View Using a Hierarchical Data Guide Obtained With GET_INDEX_DATAGUIDE	20-17
20-6	Creating a View Using a Hierarchical Data Guide Obtained With JSON_DATAGUIDE	20-18
20-7	Creating a View That Projects All Scalar Fields	20-20
20-8	Creating a View That Projects Scalar Fields Targeted By a Path Expression	20-21
20-9	Creating a View That Projects Scalar Fields Having a Given Frequency	20-22
20-10	Adding Virtual Columns That Project JSON Fields Using a Data Guide Obtained With GET_INDEX_DATAGUIDE	20-25
20-11	Adding Virtual Columns, Hidden and Visible	20-26
20-12	Projecting All Scalar Fields Not Under an Array as Virtual Columns	20-28
20-13	Projecting Scalar Fields With a Minimum Frequency as Virtual Columns	20-28
20-14	Projecting Scalar Fields With a Minimum Frequency as Hidden Virtual Columns	20-29
20-15	Dropping Virtual Columns Projected From JSON Fields	20-30
20-16	Adding Virtual Columns Automatically With Change Trigger ADD_VC	20-31
20-17	Tracing Data-Guide Updates With a User-Defined Change Trigger	20-33
20-18	Adding a 2015 Purchase-Order Document	20-36
20-19	Adding a 2016 Purchase-Order Document	20-36
20-20	Creating Multiple Data Guides With Aggregate Function JSON_DATAGUIDE	20-36
20-21	Querying a Data Guide Obtained Using JSON_DATAGUIDE	20-39
20-22	Querying a Data Guide With Index Data For Paths With Frequency at Least 80%	20-40
20-23	Flat Data Guide For Purchase Orders	20-42
20-24	Hierarchical Data Guide For Purchase Orders	20-47
21-1	Declaring an Input Value To Be JSON	21-5
21-2	Using Name–Value Pairs with JSON_OBJECT	21-7
21-3	Using Column Names with JSON_OBJECT	21-8
21-4	Using a Wildcard (*) with JSON_OBJECT	21-8
21-5	Using JSON_OBJECT With ABSENT ON NULL	21-9

21-6	Using a User-Defined Object-Type Instance with JSON_OBJECT	21-9
21-7	Using JSON_ARRAY to Construct a JSON Array	21-11
21-8	Using JSON_OBJECTAGG to Construct a JSON Object	21-12
21-9	Using JSON_ARRAYAGG to Construct a JSON Array	21-13
23-1	Constructing and Serializing an In-Memory JSON Object	23-1
23-2	Using Method GET_KEYS() to Obtain a List of Object Fields	23-2
23-3	Using Method PUT() to Update Parts of JSON Documents	23-2
24-1	A Table With GeoJSON Data	24-2
24-2	Selecting a geometry Object From a GeoJSON Feature As an SDO_GEOMETRY Instance	24-3
24-3	Retrieving Multiple geometry Objects From a GeoJSON Feature As SDO_GEOMETRY	24-4
24-4	Creating a Spatial Index For Scalar GeoJSON Data	24-5
24-5	Using GeoJSON Geometry With Spatial Operators	24-5
24-6	Creating a Materialized View Over GeoJSON Data	24-6
24-7	Creating a Spatial Index on a Materialized View Over GeoJSON Data	24-6
26-1	Creating a Bitmap Index for JSON_EXISTS	26-3
26-2	Creating a Bitmap Index for JSON_VALUE	26-3
26-3	Creating a Function-Based Index for a JSON Field: Dot Notation	26-4
26-4	Creating a Function-Based Index for a JSON Field: JSON_VALUE	26-4
26-5	Specifying NULL ON EMPTY for a JSON_VALUE Function-Based Index	26-4
26-6	Use of a JSON_VALUE Function-Based Index with a JSON_TABLE Query	26-5
26-7	JSON_EXISTS Query Targeting Field Compared to Literal Number	26-6
26-8	JSON_EXISTS Query Targeting Field Compared to Variable Value	26-7
26-9	JSON_EXISTS Query Targeting Field Cast to Number Compared to Variable Value	26-7
26-10	JSON_EXISTS Query Targeting a Conjunction of Field Comparisons	26-7
26-11	JSON_VALUE Query with Explicit RETURNING NUMBER	26-8
26-12	JSON_VALUE Query with Explicit Numerical Conversion	26-9
26-13	JSON_VALUE Query with Implicit Numerical Conversion	26-9
26-14	Creating Virtual Columns For JSON Object Fields	26-9
26-15	Creating a Composite B-tree Index For JSON Object Fields	26-9
26-16	Two Ways to Query JSON Data Indexed With a Composite Index	26-10
26-17	Creating a JSON Search Index	26-11
26-18	Execution Plan Indication that a JSON Search Index Is Used	26-11
26-19	Full-Text Query of JSON Data	26-11
26-20	Full-Text Query of JSON Data, with Escaped Search Pattern	26-12
26-21	Some Ad Hoc JSON Queries	26-12
27-1	Populating JSON Data Into the IM Column Store	27-4

List of Figures

B-1	json_basic_path_expression	B-1
B-2	json_absolute_path_expression	B-1
B-3	json_nonfunction_steps	B-1
B-4	json_object_step	B-1
B-5	json_field_name	B-1
B-6	json_array_step	B-2
B-7	json_function_step	B-2
B-8	json_item_method	B-3
B-9	json_filter_expr	B-3
B-10	json_cond	B-4
B-11	json_conjunction	B-4
B-12	json_comparison	B-4
B-13	json_relative_path-expr	B-4
B-14	json_compare_pred	B-5
B-15	json_var	B-5
B-16	json_scalar	B-5

List of Tables

5-1	JSON Object Field Syntax Examples	5-4
13-1	Compatibility of Type-Conversion Item Methods and RETURNING Types	13-12
14-1	JSON_QUERY Wrapper Clause Examples	14-5
16-1	Compatible Scalar Data Types: Converting JSON to SQL	16-4
20-1	SQL and PL/SQL Functions to Obtain a Data Guide	20-8
20-2	JSON Schema Fields (Keywords)	20-9
20-3	Oracle-Specific Data-Guide Fields	20-10
20-4	Preferred Names for Some JSON Field Columns	20-14
20-5	Parameters of a User-Defined Data-Guide Change Trigger Procedure	20-32
24-1	GeoJSON Geometry Objects Other Than Geometry Collections	24-1

Preface

This manual describes the use of JSON data that is stored in Oracle Database. It covers how to store, generate, view, manipulate, manage, search, and query it.

- [Audience](#)
Oracle Database JSON Developer's Guide is intended for developers building JSON Oracle Database applications.
- [Documentation Accessibility](#)
- [Related Documents](#)
Oracle and other resources related to this developer's guide are presented.
- [Conventions](#)
The conventions used in this document are described.
- [Code Examples](#)
The code examples in this book are for illustration only. In many cases, however, you can copy and paste parts of examples and run them in your environment.

Audience

Oracle Database JSON Developer's Guide is intended for developers building JSON Oracle Database applications.

An understanding of JSON is helpful when using this manual. Many examples provided here are in SQL or PL/SQL. A working knowledge of one of these languages is presumed.

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.

Related Documents

Oracle and other resources related to this developer's guide are presented.

- *Oracle Database New Features Guide* for information about the differences between Oracle Database Release 18c and Oracle Database Release 18c Enterprise Edition with respect to available features and options. This book also describes features new to Oracle Database Release 18c, Version 1.
- *Oracle Database Error Messages Reference*. Oracle Database error message documentation is available only as HTML. If you have access to only printed or PDF Oracle Database documentation, you can browse the error messages by range. Once you find the specific range, use the search (find) function of your Web browser to locate the specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle Database online documentation.
- *Oracle as a Document Store* for information about Simple Oracle Document Access (SODA)
- *Oracle Database Concepts*
- *Oracle Database In-Memory Guide*
- *Oracle Database SQL Language Reference*
- *Oracle Database PL/SQL Language Reference*
- *Oracle Database PL/SQL Packages and Types Reference*
- *Oracle Text Reference*
- *Oracle Text Application Developer's Guide*
- *Oracle Database Development Guide*

To download free release notes, installation documentation, white papers, or other collateral, please visit the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at OTN Registration.

If you already have a user name and password for OTN then you can go directly to the documentation section of the OTN Web site at OTN Documentation.

For additional information, see:

- ISO/IEC 13249-2:2000, Information technology - Database languages - SQL Multimedia and Application Packages - Part 2: Full-Text, International Organization For Standardization, 2000

Conventions

The conventions used in this document are described.

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.

Code Examples

The code examples in this book are for illustration only. In many cases, however, you can copy and paste parts of examples and run them in your environment.

- [Pretty Printing of JSON Data](#)
To promote readability, especially of lengthy or complex JSON data, output is sometimes shown pretty-printed (formatted) in code examples.
- [Execution Plans](#)
Some of the code examples in this book present execution plans. These are for illustration only. Running examples that are presented here in your environment is likely to result in different execution plans from those presented here.
- [Reminder About Case Sensitivity](#)
JSON is case-sensitive. SQL is case-insensitive, but names in SQL code are implicitly uppercase.

Pretty Printing of JSON Data

To promote readability, especially of lengthy or complex JSON data, output is sometimes shown pretty-printed (formatted) in code examples.

Execution Plans

Some of the code examples in this book present execution plans. These are for illustration only. Running examples that are presented here in your environment is likely to result in different execution plans from those presented here.

Reminder About Case Sensitivity

JSON is case-sensitive. SQL is case-insensitive, but names in SQL code are implicitly uppercase.

When examining the examples in this book, keep in mind the following:

- SQL is case-insensitive, but names in SQL code are implicitly uppercase, unless you enclose them in double quotation marks ("").
- JSON is case-sensitive. You must refer to SQL names in JSON code using the correct case: uppercase SQL names must be written as uppercase.

For example, if you create a table named `my_table` in SQL without using double quotation marks, then you must refer to it in JSON code as `"MY_TABLE"`.

Changes in This Release for Oracle Database JSON Developer's Guide

Changes in Oracle Database for this book are described.

Oracle Database JSON Developer's Guide was a new book in Oracle Database 12c Release 2 (12.2.0.1). Information about using JSON data in Oracle Database 12c Release 1 (12.1.0.2) is available in [Oracle XML DB Developer's Guide](#).

- [Changes in Oracle Database Release 19c, Version 19.1, for JSON Developer's Guide](#)
The changes in JSON support and in *Oracle Database JSON Developer's Guide* for Oracle Database Release 19c, Version 19.1, are described.
- [Changes in Oracle Database Release 18c, Version 18.1, for JSON Developer's Guide](#)
The changes in JSON support and in *Oracle Database JSON Developer's Guide* for Oracle Database Release 18c, Version 18.1, are described.

Changes in Oracle Database Release 19c, Version 19.1, for JSON Developer's Guide

The changes in JSON support and in *Oracle Database JSON Developer's Guide* for Oracle Database Release 19c, Version 19.1, are described.

- [New Features](#)
The following features are new in this release.

New Features

The following features are new in this release.

- [JSON Materialized View Support](#)
Performance enhancement: If you create a refresh-on-statement materialized view over `json_table` and some other conditions apply then a query that matches the query defining the view can be rewritten to a materialized-view table access. You can use this feature instead of creating multiple functional indexes.
- [SQL Function JSON_MERGEPATCH: Declarative Update of JSON Documents](#)
You can now update a JSON document declaratively, using new SQL function `json_mergepatch`. You can apply one or more changes to multiple documents using a single statement. This feature improves the flexibility of JSON update operations.

- [New SQL/JSON Function JSON_SERIALIZE and JSON Data Guide Support for GeoJSON Data](#)
You can use new SQL function `json_serialize` to serialize JSON data to text or to UTF-encoded BLOB data. SQL aggregate function `json_dataguide` can now detect GeoJSON geographic data in your documents. You can use it to create a view that projects such data as SQL data type `SDO_GEOMETRY`.
- [Syntax Simplifications](#)
Syntax simplifications are offered for SQL/JSON path expressions and SQL/JSON generation with function `json_object`. A new SQL query clause, `NESTED`, provides a simple alternative to using `json_table` with `LEFT OUTER JOIN`.
- [Mapping of JSON Data To and From SQL Object Types](#)
This feature enables the mapping of JSON data to and from user-defined SQL object types and collections. You can convert JSON data to an instance of a SQL object type using SQL/JSON function `json_value`. In the opposite direction, you can generate JSON data from an instance of a SQL object type using SQL/JSON function `json_object` or `json_array`.

JSON Materialized View Support

Performance enhancement: If you create a refresh-on-statement materialized view over `json_table` and some other conditions apply then a query that matches the query defining the view can be rewritten to a materialized-view table access. You can use this feature instead of creating multiple functional indexes.

SQL Function JSON_MERGEPATCH: Declarative Update of JSON Documents

You can now update a JSON document declaratively, using new SQL function `json_mergepatch`. You can apply one or more changes to multiple documents using a single statement. This feature improves the flexibility of JSON update operations.

Related Topics

- [Updating a JSON Document with JSON Merge Patch](#)
You can use Oracle SQL function `json_mergepatch` to update specific portions of a JSON document. You pass it a JSON Merge Patch document, which specifies the changes to make to a specified JSON document. JSON Merge Patch is an IETF standard.

New SQL/JSON Function JSON_SERIALIZE and JSON Data Guide Support for GeoJSON Data

You can use new SQL function `json_serialize` to serialize JSON data to text or to UTF-encoded BLOB data. SQL aggregate function `json_dataguide` can now detect GeoJSON geographic data in your documents. You can use it to create a view that projects such data as SQL data type `SDO_GEOMETRY`.

Related Topics

- [Oracle SQL Function JSON_SERIALIZE](#)
Oracle SQL function `json_serialize` takes JSON data (of any SQL data type, `VARCHAR2`, `CLOB`, or `BLOB`) as input and returns a textual representation of it (as `VARCHAR2`, `CLOB`, or `BLOB` data). `VARCHAR2(4000)` is the default return type.

- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.

Syntax Simplifications

Syntax simplifications are offered for SQL/JSON path expressions and SQL/JSON generation with function `json_object`. A new SQL query clause, `NESTED`, provides a simple alternative to using `json_table` with `LEFT OUTER JOIN`.

Related Topics

- [Basic SQL/JSON Path Expression Syntax](#)
The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (\$) followed by zero or more object, array, and descendant steps, each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.
- [Overview of SQL/JSON Generation Functions](#)
You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.
- [SQL/JSON Function JSON_OBJECT](#)
SQL/JSON function `json_object` constructs JSON objects from the results of evaluating its argument SQL expressions.
- [SQL NESTED Clause Instead of JSON_TABLE](#)
In a `SELECT` clause you can often use a `NESTED` clause instead of SQL/JSON function `json_table`. This can mean a simpler query expression. It also has the advantage of including rows with non-NULL relational columns when the JSON column is NULL.
- [Diagrams for Basic SQL/JSON Path Expression Syntax](#)
Syntax diagrams and corresponding Backus-Naur Form (BNF) syntax descriptions are presented for the basic SQL/JSON path expression syntax.

Mapping of JSON Data To and From SQL Object Types

This feature enables the mapping of JSON data to and from user-defined SQL object types and collections. You can convert JSON data to an instance of a SQL object type using SQL/JSON function `json_value`. In the opposite direction, you can generate JSON data from an instance of a SQL object type using SQL/JSON function `json_object` or `json_array`.

Related Topics

- [SQL/JSON Function JSON_VALUE](#)
SQL/JSON function `json_value` selects JSON data and returns a SQL scalar or an instance of a user-defined SQL object type or SQL collection type (varray, nested table).
- [Using JSON_VALUE To Instantiate a User-Defined Object Type Instance](#)
You can use SQL/JSON function `json_value` to instantiate an instance of a user-defined SQL object type or collection type. You do this by targeting a JSON object or array in the path expression and specifying the object or collection type, respectively, in the `RETURNING` clause.

- [ON MISMATCH Clause for JSON_VALUE](#)
When the `RETURNING` clause specifies a user-defined object-type or collection-type instance, function `json_value` accepts an optional `ON MISMATCH` clause, which specifies handling to use when a targeted JSON value does not match the specified SQL return value. This clause and its default behavior (no `ON MISMATCH` clause) are described here.
- [SQL/JSON Function JSON_OBJECT](#)
SQL/JSON function `json_object` constructs JSON objects from the results of evaluating its argument SQL expressions.
- [SQL/JSON Function JSON_ARRAY](#)
SQL/JSON function `json_array` constructs a JSON array from the results of evaluating its argument SQL expressions.

Changes in Oracle Database Release 18c, Version 18.1, for JSON Developer's Guide

The changes in JSON support and in *Oracle Database JSON Developer's Guide* for Oracle Database Release 18c, Version 18.1, are described.

- [New Features](#)
The following features are new in this release.
- [Deprecated Features](#)
The following features are deprecated in this release, and may be desupported in a future release.

New Features

The following features are new in this release.

- [SQL Enhancements](#)
Multiple SQL enhancements are provided for storing and querying JSON data in the database.
- [Sharding Support](#)
Sharding support for JSON data has been enhanced.
- [Performance Improvements for LOB Storage](#)
Performance has been improved for JSON data that uses LOB storage. This particularly affects the common use case of reading and writing an entire JSON document stored as a LOB.
- [JSON Search Index Support for Longer Field Names](#)
The maximum JSON field name length supported for a JSON search index is now 255 bytes (formerly it was 64 bytes).

SQL Enhancements

Multiple SQL enhancements are provided for storing and querying JSON data in the database.

You Can Specify That a SQL Expression Returns JSON Data

You can apply SQL function `treat`, with keywords `AS JSON`, to a SQL expression to specify that the expression returns JSON data. This is useful in situations where Oracle cannot determine that the result is JSON data. It is also useful in situations where you want to force some text to be interpreted as JSON data. For example, you can use it to tell Oracle to interpret a `VARCHAR2` value of `{ }` not as a string but as an empty JSON object.

See Also:

Oracle Database SQL Language Reference

LOB Results for `JSON_VALUE`, `JSON_QUERY`, and `JSON_TABLE`

SQL/JSON function `json_value` can now return a `CLOB` instance.

SQL/JSON function `json_query` can now return a `CLOB` or `BLOB` instance. A `BLOB` result is in the `AL32UTF8` character set.

As before, the data type of a `json_table` column depends on whether the column is specified as `FORMAT JSON`. If it is, the `json_query` return types are supported; if it is not, the `json_value` return types are supported.

Previously:

- `json_value` supported only `VARCHAR2`, `NUMBER`, `DATE`, `TIMESTAMP`, `TIMESTAMP WITH TIME ZONE`, and `SDO_GEOMETRY` as return types.
- `json_query` supported only `VARCHAR2` as return type.

See Also:

Oracle Database SQL Language Reference

SQL/JSON Generation Function Enhancements

- Optional keyword `STRICT` has been added. When present, any inputs that are declared or otherwise expected to be JSON data are checked to ensure that they are, in fact, well-formed JSON data. If any of them is not then an error is raised.

(You can declare input to be JSON by using keywords `FORMAT JSON` or by using SQL function `treat` with keywords `AS JSON`. Input is expected to be JSON data if it is from a table column that has an `is json` constraint or it is the output of another JSON generation function.)

- The SQL/JSON generation functions (`json_object`, `json_array`, `json_objectagg`, and `json_arrayagg`) can now return a CLOB or BLOB instance. A BLOB result is in the AL32UTF8 character set.
- The input expression to a generation function can now be of any of these SQL data types (in addition to NUMBER, VARCHAR2, DATE and TIMESTAMP):
 - BINARY_DOUBLE
 - BINARY_FLOAT
 - CLOB
 - BLOB
 - NVARCHAR2
 - RAW
 - TIMESTAMP WITH TIME ZONE
 - TIMESTAMP WITH LOCAL TIME ZONE
 - INTERVAL YEAR TO MONTH
 - INTERVAL DAY TO SECOND

How the result of generation is rendered in JSON depends on the data type of the input.

See Also:

- [Handling of Input Values For SQL/JSON Generation Functions](#)
- *Oracle Database SQL Language Reference*

Item Methods for SQL/JSON Path Expressions

- New item methods have been added: `numberOnly()`, `stringOnly()`, `boolean()`, `booleanOnly()`, `size()`, and `type()`.

The data-type conversion methods with “only” in their name are the same as the corresponding methods with names without “only”, except that the former convert *only* JSON values that are of the given type (e.g., `number`) to the related SQL data type (e.g. `NUMBER`). The methods without “only” in the name allow conversion, when possible, of *any* JSON value to the given SQL data type.

- All item methods can now be used in path expressions for SQL/JSON functions `json_value`, `json_query`, and `json_table` (columns). In Oracle Database 12c Release 2 (12.2.0.1), item methods could only be used in path expressions for SQL/JSON condition `json_exists`.
- Item method `string()` can now return CLOB or BLOB (in addition to VARCHAR2). The default is `VARCHAR2(4000)`.

 **See Also:**

[SQL/JSON Path Expression Item Methods](#)

JSON Data Guide Enhancements

- Oracle SQL aggregate function `json_dataguide` now accepts optional formatting arguments:
 - You can return a *hierarchical* dataguide, by specifying keyword **FORMAT** with argument `DBMS_JSON.FORMAT_HIERARCHICAL`.
 - You can specify *pretty-printing* (indentation to improve readability) of the returned dataguide, by specifying keyword **PRETTY** with argument `DBMS_JSON.PRETTY`.
- New data-dictionary views are available, to show you the dataguide information recorded for individual JSON object *fields* in a dataguide-enabled search index: `USER_JSON_DATAGUIDE_FIELDS`, `ALL_JSON_DATAGUIDE_FIELDS`, `DBA_JSON_DATAGUIDE_FIELDS`. (These are in addition to views `USER_JSON_DATAGUIDES`, `ALL_JSON_DATAGUIDES`, and `DBA_JSON_DATAGUIDES`, which list the tables that have JSON columns with a dataguide-enabled search index.)
- For JSON documents that contain an array of scalar values, a dataguide now records the scalar types as well as the type `ARRAY`. The recorded path for the scalar values is indicated in a flat data guide by appending `[*]` to the path recorded for the array itself.

 **See Also:**

Oracle Database SQL Language Reference

Data-Dictionary Views That Record the Presence of JSON Columns

Data-dictionary views `USER_JSON_COLUMNS`, `ALL_JSON_COLUMNS`, and `DBA_JSON_COLUMNS` now list the *views*, as well as the tables, that have columns with JSON data.

 **See Also:**

Oracle Database Reference for information about `ALL_JSON_COLUMNS` and the related data-dictionary views

SQL/JSON Function `JSON_TABLE` Syntax

The syntax of `json_table` has been enhanced by making it simpler for some common use cases:

- You can now use simple dot-notation syntax in place of a path expression.

- If a column is the projection of a JSON object field, and if you want the column to have the same name as the field, then you need not provide a path expression to that object — the path is inferred from the column name.



See Also:

Oracle Database SQL Language Reference

ON STATEMENT Support For JSON_TABLE Materialized Views

You can now use keywords `ON STATEMENT` when creating a materialized view using a `json_table` query. Using `ON STATEMENT` instead of `ON COMMIT` means that the view is automatically synchronized for each DML statement against the base table.

New SQL Function TO_UTC_TIMESTAMP_TZ

SQL function `to_utc_timestamp_tz` takes as input an ISO 8601 date format string and returns an instance of SQL data type `TIMESTAMP WITH TIMEZONE`. It normalizes the input to UTC time (Coordinated Universal Time, formerly Greenwich Mean Time). Unlike SQL function `to_timestamp_tz`, the new function assumes that the input string uses the ISO 8601 date format, defaulting the time zone to UTC 0.

A typical use of this function would be to provide its output to SQL function `sys_extract_utc`, obtaining a UTC time that is then passed as a SQL bind variable to SQL/JSON condition `json_exists`, to perform a time-stamp range comparison.



See Also:

ISO 8601 for information about the ISO date formats

New Oracle SQL Condition JSON_EQUAL

Oracle SQL condition `json_equal` compares two JSON values and returns true if they are equal, false otherwise. For this comparison, insignificant whitespace and insignificant object member order are ignored. For example, JSON objects are equal if they have the same members, regardless of their order. However, if either of two compared objects has one or more duplicate fields then the value returned by `json_equal` is unspecified.

Related Topics

- [Determining Whether a Column Necessarily Contains JSON Data](#)
How can you tell whether a given column of a table or view is well-formed JSON data? Whenever this is the case, the column is listed in the following static data dictionary views: `DBA_JSON_COLUMNS`, `USER_JSON_COLUMNS`, and `ALL_JSON_COLUMNS`.
- [Basic SQL/JSON Path Expression Syntax](#)
The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (\$) followed by zero or more object, array, and descendant steps, each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.

- [Data-Dictionary Views For Persistent Data-Guide Information](#)
You can query static data-dictionary views to see which tables have JSON columns with data guide-enabled JSON search indexes and to extract JSON object field information that is recorded in dataguide-enabled JSON search indexes.

Sharding Support

Sharding support for JSON data has been enhanced.

The following are now supported:

- JSON search index on sharded tables, whether user-managed or system-managed.
- Cross-shard (and in-shard) queries for JSON data returned as CLOB or BLOB from SQL and PL/SQL. This includes projection of JSON data to CLOB and BLOB columns using SQL/JSON function `json_table`.
- Evaluation on individual shards of cross-shard queries that use SQL/JSON functions and conditions. This includes the use of the following:
 - Function `json_value` and aggregate function `json_dataguide` in a SELECT clause
 - Function `json_table` in a FROM clause
 - Function `json_value`, and conditions `json_exists` and `json_textcontains`, in a WHERE clause.
- Automatic relocation of base-table and index storage-table partitions during chunk migration.

See Also:

Using Oracle Sharding

Performance Improvements for LOB Storage

Performance has been improved for JSON data that uses LOB storage. This particularly affects the common use case of reading and writing an entire JSON document stored as a LOB.

JSON Search Index Support for Longer Field Names

The maximum JSON field name length supported for a JSON search index is now 255 bytes (formerly it was 64 bytes).

 **Note:**

You **must rebuild** any JSON search indexes and Oracle Text indexes created prior to Oracle Database 18c if they index JSON data that contains object fields with names longer than 64 bytes. See *Oracle Database Upgrade Guide* for more information.

Related Topics

- [Oracle Database JSON Restrictions](#)
The restrictions associated with Oracle support of JSON data in Oracle Database are listed here.

Deprecated Features

The following features are deprecated in this release, and may be desupported in a future release.

- SQL/JSON functions returning a Boolean JSON value as a number (zero or one). Return the value as `VARCHAR2` or (in PL/SQL only) as `BOOLEAN`; do not return it as `NUMBER`. If you really need a SQL numeric value then you can use `SQL DECODE` or `CASE WHEN` to obtain zero or one from a `VARCHAR2` value.

Part I

Introduction to JSON Data and Oracle Database

Get started understanding JSON data and how you can use SQL and PL/SQL with JSON data stored in Oracle Database.

Schemaless development based on persisting application data in the form of JSON documents lets you quickly react to changing application requirements. You can change and redeploy your application without needing to change the storage schemas it uses.

SQL and relational databases provide flexible support for complex data analysis and reporting, as well as rock-solid data protection and access control. This is typically *not* the case for NoSQL databases, which have often been associated with schemaless development with JSON in the past.

Oracle Database provides all of the benefits of SQL and relational databases to JSON data, which you store and manipulate in the same ways and with the same confidence as any other type of database data.

- [JSON in Oracle Database](#)
Oracle Database supports JavaScript Object Notation (JSON) data natively with relational database features, including transactions, indexing, declarative querying, and views.
- [JSON Data](#)
JavaScript Object Notation (JSON) is defined in standards ECMA-404 (JSON Data Interchange Format) and ECMA-262 (ECMAScript Language Specification, third edition). The JavaScript dialect of ECMAScript is a general programming language used widely in web browsers and web servers.

1

JSON in Oracle Database

Oracle Database supports JavaScript Object Notation (JSON) data natively with relational database features, including transactions, indexing, declarative querying, and views.

This documentation covers the use of database languages and features to work with JSON data that is stored in Oracle Database. In particular, it covers how to use SQL and PL/SQL with JSON data.

Note:

Oracle also provides a family of **Simple Oracle Document Access (SODA)** APIs for access to JSON data stored in the database. SODA is designed for schemaless application development without knowledge of relational database features or languages such as SQL and PL/SQL. It lets you create and store collections of documents in Oracle Database, retrieve them, and query them, without needing to know how the documents are stored in the database.

There are several implementations of SODA:

- SODA for REST — Representational state transfer (REST) requests perform collection and document operations, using any language capable of making HTTP calls.
- SODA for Java — Java classes and interfaces represent databases, collections, and documents.
- SODA for PL/SQL — PL/SQL object types represent collections and documents.
- SODA for C — Oracle Call Interface (OCI) handles represent collections and documents.

For information about SODA see [Oracle as a Document Store](#).

- [Overview of JSON in Oracle Database](#)
Oracle Database supports JSON natively with relational database features, including transactions, indexing, declarative querying, and views. Unlike relational data, JSON data can be stored in the database, indexed, and queried without any need for a schema that defines the data.
- [Getting Started Using JSON with Oracle Database](#)
In general, you will perform the following tasks when working with JSON data in Oracle Database: (1) create a JSON column with an `is json` check constraint, (2) insert JSON data into the column, and (3) query the JSON data.

- [Oracle Database Support for JSON](#)
Oracle Database support for JavaScript Object Notation (JSON) is designed to provide the best fit between the worlds of relational storage and querying JSON data, allowing relational and JSON queries to work well together. Oracle SQL/JSON support is closely aligned with the JSON support in the SQL Standard.

1.1 Overview of JSON in Oracle Database

Oracle Database supports JSON natively with relational database features, including transactions, indexing, declarative querying, and views. Unlike relational data, JSON data can be stored in the database, indexed, and queried without any need for a schema that defines the data.

(The JSON data is schemaless, even though a *database schema* is used to define the table and column in which it is stored. Nothing in that schema specifies the structure of the JSON data itself.)

JSON data has often been stored in NoSQL databases such as Oracle NoSQL Database and Oracle Berkeley DB. These allow for storage and retrieval of data that is not based on any schema, but they do not offer the rigorous consistency models of relational databases.

To compensate for this shortcoming, a relational database is sometimes used in parallel with a NoSQL database. Applications using JSON data stored in the NoSQL database must then ensure data integrity themselves.

Native support for JSON by Oracle Database obviates such workarounds. It provides *all* of the benefits of relational database features for use with JSON, including transactions, indexing, declarative querying, and views.

Database queries with Structured Query Language (SQL) are declarative. With Oracle Database you can use SQL to join JSON data with relational data. And you can project JSON data relationally, making it available for relational processes and tools. You can also query, from within the database, JSON data that is stored outside Oracle Database in an external table.

You can access JSON data stored in the database the same way you access other database data, including using Oracle Call Interface (OCI), and Java Database Connectivity (JDBC).

The JSON-language standard defines JSON data in a *textual* way: it is composed of Unicode characters in a standard syntax. In Oracle Database, you can store JSON data textually using the common SQL data types `VARCHAR2`, `CLOB`, and `BLOB`, as unparsed character data.

Oracle recommends that you *always* use an `is json` check constraint to ensure that column values are valid JSON instances. See [Example 4-1](#).

JSON Columns in Database Tables

Oracle Database places no restrictions on the tables that can be used to store JSON documents. A column containing JSON documents can coexist with any other kind of database data. A table can also have multiple columns that contain JSON documents.

When using Oracle Database as a JSON document store, your tables that contain JSON columns typically also have a few non-JSON housekeeping columns. These typically track metadata about the JSON documents.

If you are using JSON to add flexibility to a primarily relational application then some of your tables likely also have a column for JSON documents, which you use to manage the application data that does not map directly to your relational model.

By definition, textual JSON data is encoded using a Unicode encoding, either UTF-8 or UTF-16. You can use textual data that is stored in a non-Unicode character set as if it were JSON data, but in that case Oracle Database automatically converts the character set to UTF-8 when processing the data.

Use SQL With JSON Data

In SQL, you can access JSON data stored in Oracle Database using either specialized functions and conditions or a simple dot notation. Most of the SQL functions and conditions belong to the SQL/JSON standard, but a few are Oracle-specific.

- SQL/JSON query functions `json_value`, `json_query`, and `json_table`.
These evaluate SQL/JSON path expressions against JSON data to produce SQL values.

- Oracle SQL condition `json_textcontains` and SQL/JSON conditions `json_exists`, `is json`, and `is not json`.

Condition `json_exists` checks for the existence of given JSON data; `json_textcontains` provides full-text querying of JSON data; and `is json` and `is not json` check whether given JSON data is well-formed.

`json_exists` and `json_textcontains` check the data that matches a SQL/JSON path expression.

- A simple *dot notation* that acts similar to a combination of query functions `json_value` and `json_query`.

This resembles a SQL object access expression, that is, attribute dot notation for an abstract data type (ADT). This is the *easiest* way to query JSON data in the database.

- SQL/JSON *generation* functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg`.

These gather SQL data to produce JSON object and array data (as a SQL value).

- Oracle SQL condition `json_equal`, which tests whether two JSON values are the same.

- Oracle SQL aggregate function `json_dataguide`.

This produces JSON data that is a *data guide*, which you can use to discover information about the structure and content of other JSON data in the database.

As a simple illustration of querying, here is a dot-notation query of the documents stored in JSON column `po_document` of table `j_purchaseorder` (aliased here as `po`). It obtains all purchase-order requestors (JSON field `Requestor`).

```
SELECT po.po_document.Requestor FROM j_purchaseorder po;
```

Use PL/SQL With JSON Data

You can manipulate JSON data within PL/SQL code using SQL code or PL/SQL object types for JSON. (You cannot use an empty JSON field name in any SQL code that you use in PL/SQL.)

The following SQL functions and conditions are also available as built-in PL/SQL functions: `json_value`, `json_query`, `json_object`, `json_array`, `json_exists`, `is json`, `is not json`, and `json_equal`.

Unlike the case for Oracle SQL, which has no `BOOLEAN` data type, in *PL/SQL*:

- `json_exists`, `is json`, `is not json`, and `json_equal` are Boolean functions.
- `json_value` can return a `BOOLEAN` value.

There are also PL/SQL object types for JSON, which you can use for fine-grained construction and manipulation of In-Memory JSON data. You can introspect it, modify it, and serialize it back to textual JSON data.

Related Topics

- [Simple Dot-Notation Access to JSON Data](#)
Dot notation is designed for easy, general use and common use cases. Queries of JSON data that use dot-notation syntax return JSON values whenever possible.
- [Overview of SQL/JSON Path Expressions](#)
Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.
- [JSON Data Guide](#)
A JSON data guide lets you discover information about the structure and content of JSON documents stored in Oracle Database.
- [Oracle Database Support for JSON](#)
Oracle Database support for JavaScript Object Notation (JSON) is designed to provide the best fit between the worlds of relational storage and querying JSON data, allowing relational and JSON queries to work well together. Oracle SQL/JSON support is closely aligned with the JSON support in the SQL Standard.
- [Character Sets and Character Encoding for JSON Data](#)
JSON data always uses the Unicode character set. In this respect, JSON data is simpler to use than XML data. This is an important part of the JSON Data Interchange Format (RFC 4627). For JSON data processed by Oracle Database, any needed character-set conversions are performed automatically.
- [PL/SQL Object Types for JSON](#)
You can use PL/SQL object types for JSON to read and write multiple fields of a JSON document. This can increase performance, in particular by avoiding multiple parses and serializations of the data.
- [Overview of Storing and Managing JSON Data](#)
This overview describes: (1) data types for JSON columns, (2) LOB storage considerations for JSON data, and (3) ensuring that JSON columns contain well-formed JSON data.

1.2 Getting Started Using JSON with Oracle Database

In general, you will perform the following tasks when working with JSON data in Oracle Database: (1) create a JSON column with an `is json` check constraint, (2) insert JSON data into the column, and (3) query the JSON data.

1. Create a table with a primary-key column and a JSON column, and add an `is json` *check constraint* to ensure that the JSON column contains only well-formed JSON data.

The following statement creates table `j_purchaseorder` with primary key `id` and with JSON column `po_document` (see also [Example 4-1](#)).

```
CREATE TABLE j_purchaseorder
  (id          VARCHAR2 (32) NOT NULL PRIMARY KEY,
   date_loaded  TIMESTAMP (6) WITH TIME ZONE,
   po_document  VARCHAR2 (32767)
   CONSTRAINT ensure_json CHECK (po_document IS JSON));
```

2. Insert JSON data into the JSON column, using any of the methods available for Oracle Database.

The following statement uses a SQL `INSERT` statement to insert some simple JSON data into the third column of table `j_purchaseorder` (which is column `po_document` — see previous). Some of the JSON data is elided here (...). See [Example 4-2](#) for these details.

```
INSERT INTO j_purchaseorder
  VALUES (SYS_GUID(),
          to_date('30-DEC-2014'),
          '{"PONumber"           : 1600,
           "Reference"          : "ABULL-20140421",
           "Requestor"         : "Alexis Bull",
           "User"              : "ABULL",
           "CostCenter"        : "A50",
           "ShippingInstructions" : {...},
           "Special Instructions" : null,
           "AllowPartialShipment" : true,
           "LineItems"         : [...]}');
```

3. Query the JSON data. The return value is always a `VARCHAR2` instance that represents a JSON value. Here are some simple examples.

The following query extracts, from each document in JSON column `po_document`, a *scalar* value, the JSON number that is the value of field `PONumber` for the objects in JSON column `po_document` (see also [Example 12-1](#)):

```
SELECT po.po_document.PONumber FROM j_purchaseorder po;
```

The following query extracts, from each document, an *array* of JSON phone objects, which is the value of field `Phone` of the object that is the value of field `ShippingInstructions` (see also [Example 12-2](#)):

```
SELECT po.po_document.ShippingInstructions.Phone FROM
j_purchaseorder po;
```

The following query extracts, from each document, *multiple* values as an array: the value of field `type` for each object in array `Phone`. The returned array is not part of the stored data but is constructed automatically by the query. (The order of the array elements is unspecified.)

```
SELECT po.po_document.ShippingInstructions.Phone.type FROM
j_purchaseorder po;
```


Related Topics

- [Creating a Table With a JSON Column](#)
You can create a table that has JSON columns. You use SQL condition `is json` as a check constraint to ensure that data inserted into a column is (well-formed) JSON data. Oracle recommends that you *always* use an `is_json` check constraint when you create a column intended for JSON data.
- [Simple Dot-Notation Access to JSON Data](#)
Dot notation is designed for easy, general use and common use cases. Queries of JSON data that use dot-notation syntax return JSON values whenever possible.

1.3 Oracle Database Support for JSON

Oracle Database support for JavaScript Object Notation (JSON) is designed to provide the best fit between the worlds of relational storage and querying JSON data, allowing relational and JSON queries to work well together. Oracle SQL/JSON support is closely aligned with the JSON support in the SQL Standard.

See Also:

- *ISO/IEC 9075-2:2016, Information technology—Database languages—SQL—Part 2: Foundation (SQL/Foundation)*
- ISO/IEC TR 19075-6
- *Oracle Database SQL Language Reference*
- JSON.org
- ECMA International

2

JSON Data

JavaScript Object Notation (JSON) is defined in standards ECMA-404 (JSON Data Interchange Format) and ECMA-262 (ECMAScript Language Specification, third edition). The JavaScript dialect of ECMAScript is a general programming language used widely in web browsers and web servers.

- [Overview of JSON](#)

JavaScript Object Notation (JSON) is defined in standards ECMA-404 (JSON Data Interchange Format) and ECMA-262 (ECMAScript Language Specification, third edition). The JavaScript dialect of ECMAScript is a general programming language used widely in web browsers and web servers.

- [JSON Syntax and the Data It Represents](#)

JSON (and JavaScript) values, scalars, objects, and arrays are described.

- [JSON Compared with XML](#)

Both JSON and XML (Extensible Markup Language) are commonly used as data-interchange languages. Their main differences are listed here.

2.1 Overview of JSON

JavaScript Object Notation (JSON) is defined in standards ECMA-404 (JSON Data Interchange Format) and ECMA-262 (ECMAScript Language Specification, third edition). The JavaScript dialect of ECMAScript is a general programming language used widely in web browsers and web servers.

JSON is almost a subset of the object literal notation of JavaScript.¹ Because it can be used to represent JavaScript object literals, JSON commonly serves as a data-interchange language. In this it has much in common with XML.

Because it is (almost a subset of) JavaScript notation, JSON can often be used in JavaScript programs without any need for parsing or serializing. It is a text-based way of representing JavaScript object literals, arrays, and scalar data.

Although it was defined in the context of JavaScript, JSON is in fact a language-independent data format. A variety of programming languages can parse and generate JSON data.

JSON is relatively easy for humans to read and write, and easy for software to parse and generate. It is often used for serializing structured data and exchanging it over a network, typically between a server and web applications.

¹ JSON differs from JavaScript notation in this respect: JSON allows unescaped Unicode characters U+2028 (LINE SEPARATOR) and U+2029 (PARAGRAPH SEPARATOR) in strings. JavaScript notation requires control characters such as these to be escaped in strings. This difference can be important when generating JSONP (JSON with padding) data.

 **See Also:**

- ECMA 404 and IETF RFC4627 for the definition of the JSON Data Interchange Format
- ECMA 262 for the ECMAScript Language Specification
- JSON.org for information about JSON

2.2 JSON Syntax and the Data It Represents

JSON (and JavaScript) values, scalars, objects, and arrays are described.

A JSON **value** is one of the following: object, array, number, string, Boolean (`true` or `false`), or `null`. All values except objects and arrays are **scalar**.

 **Note:**

A JSON value of `null` is a *value* as far as SQL is concerned. It is not `NULL`, which in SQL represents the *absence* of a value (missing, unknown, or inapplicable data). In particular, SQL condition `IS NULL` returns false for a JSON `null` value, and SQL condition `IS NOT NULL` returns true.

A **JavaScript object** is an associative array, or dictionary, of zero or more pairs of **property** names and associated JSON values.² A **JSON object** is a **JavaScript object literal**.³ It is written as such a property list enclosed in braces (`{, }`), with name–value pairs separated by commas (`,`), and with the name and value of each pair separated by a colon (`:`). (Whitespace before or after the comma or colon is optional and insignificant.)

In JSON each property name and each string value *must* be enclosed in double quotation marks (`"`). In JavaScript notation, a property name used in an object literal can be, but need not be, enclosed in double quotation marks. It can also be enclosed in single quotation marks (`'`).

As a result of this difference, in practice, data that is represented using unquoted or single-quoted property names is sometimes referred to loosely as being represented in JSON, and some implementations of JSON, including the Oracle Database implementation, support the *lax syntax that allows the use of unquoted and single-quoted property names*.

A string in JSON is composed of Unicode characters, with backslash (`\`) escaping. A JSON number (numeral) is represented in decimal notation, possibly signed and possibly including a decimal exponent.

An object property is typically called a **field**. It is sometimes called a **key**, but this documentation generally uses “field” to avoid confusion with other uses here of the

² JavaScript objects are thus similar to hash tables in C and C++, HashMaps in Java, associative arrays in PHP, dictionaries in Python, and hashes in Perl and Ruby.

³ An object is created in JavaScript using either constructor `Object` or object literal syntax: `{...}`.

word “key”. An object property name–value pair is often called an object **member** (but sometimes **member** can mean just the property). Order is not significant among object members.

 **Note:**

- A JSON field name can be *empty* (written "").⁴
- Each field name in a given JSON object is not necessarily unique; the same field name can be repeated. The SQL/JSON *path evaluation* that Oracle Database employs always uses only one of the object members that have a given field name; any *other members with the same name are ignored*. It is unspecified which of multiple such members is used.

See also [Unique Versus Duplicate Fields in JSON Objects](#).

A **JavaScript array** has zero or more elements. A **JSON array** is represented by brackets ([,]) surrounding the representations of the array **elements** (also called **items**), which are separated by commas (,), and each of which is an object, an array, or a scalar value. Array *element order is significant*. (Whitespace before or after a bracket or comma is optional and insignificant.)

Example 2-1 A JSON Object (Representation of a JavaScript Object Literal)

This example shows a JSON object that represents a purchase order, with top-level field names PONumber, Reference, Requestor, User, Costcenter, ShippingInstruction, Special Instructions, AllowPartialShipment and LineItems.

```
{ "PONumber"           : 1600,
  "Reference"         : "ABULL-20140421",
  "Requestor"        : "Alexis Bull",
  "User"              : "ABULL",
  "CostCenter"       : "A50",
  "ShippingInstructions" : { "name"      : "Alexis Bull",
                           "Address" : { "street" : "200 Sporting Green",
                                         "city"   : "South San Francisco",
                                         "state"  : "CA",
                                         "zipCode" : 99236,
                                         "country" : "United States of America" },
                           "Phone"  : [ { "type" : "Office", "number" :
"909-555-7307" },
                                         { "type" : "Mobile", "number" :
"415-555-1234" } ] },
  "Special Instructions" : null,
  "AllowPartialShipment" : false,
  "LineItems"          : [ { "ItemNumber" : 1,
                           "Part"       : { "Description" : "One Magic Christmas",
                                             "UnitPrice"   : 19.95,
                                             "UPCCode"    : 13131092899 },
                           "Quantity"   : 9.0 } ],
}
```

⁴ In a few contexts an empty field name cannot be used with Oracle Database. Wherever it can be used, the name *must* be wrapped with double quotation marks.

```
{ "ItemNumber" : 2,
  "Part"       : { "Description" : "Lethal Weapon",
                  "UnitPrice"   : 19.95,
                  "UPCCode"     : 85391628927 },
  "Quantity"  : 5.0 } ] }
```

- Most of the fields here have string values. For example: field `User` has value `"ABULL"`.
- Fields `PONumber` and `zipCode` have numeric values: 1600 and 99236.
- Field `ShippingInstructions` has an object as its value. This object has three members, with fields `name`, `Address`, and `Phone`. Field `name` has a string value (`"Alexis Bull"`).
- The value of field `Address` is an object with fields `street`, `city`, `state`, `zipCode`, and `country`. Field `zipCode` has a numeric value; the others have string values.
- Field `Phone` has an array as value. This array has two elements, each of which is an object. Each of these objects has two members: fields `type` and `number` with their values.
- Field `Special Instructions` has a null value.
- Field `AllowPartialShipment` has the Boolean value `false`.
- Field `LineItems` has an array as value. This array has two elements, each of which is an object. Each of these objects has three members, with fields `ItemNumber`, `Part`, and `Quantity`.
- Fields `ItemNumber` and `Quantity` have numeric values. Field `Part` has an object as value, with fields `Description`, `UnitPrice`, and `UPCCode`. Field `Description` has a string value. Fields `UnitPrice` and `UPCCode` have numeric values.

Related Topics

- [About Strict and Lax JSON Syntax](#)
The Oracle default syntax for JSON is lax. In particular: it reflects the JavaScript syntax for object fields; the Boolean and null values are not case-sensitive; and it is more permissive with respect to numerals, whitespace, and escaping of Unicode characters.



See Also:

[Example 4-2](#)

2.3 JSON Compared with XML

Both JSON and XML (Extensible Markup Language) are commonly used as data-interchange languages. Their main differences are listed here.

JSON is most useful with simple, structured data. XML is useful for both structured and semi-structured data. JSON is generally data-centric, not document-centric; XML can be either. JSON is not a markup language; it is designed only for data

representation. XML is both a document markup language and a data representation language.

- JSON data types are few and predefined. XML data can be either typeless or based on an XML schema or a document type definition (DTD).
- JSON has simple structure-defining and document-combining constructs: it lacks attributes, namespaces, inheritance, and substitution.
- The order of the members of a JavaScript object literal is insignificant. In general, order matters within an XML document.
- JSON lacks an equivalent of XML text nodes (XPath node test `text()`). In particular, this means that there is no mixed content (which is another way of saying that JSON is not a markup language).
- JSON has no *date* data type (unlike both XML and JavaScript). A date is represented in JSON using the available data types, such as *string*. There are some de facto standards for converting between dates and JSON strings. But programs using JSON must, one way or another, deal with date representation conversion.

Because of its simple definition and features, JSON data is generally easier to generate, parse, and process than XML data. Use cases that involve combining different data sources generally lend themselves well to the use of XML, because it offers namespaces and other constructs facilitating modularity and inheritance.

Part II

Store and Manage JSON Data

This part covers creating JSON columns in a database table, partitioning such tables, replicating them using Oracle GoldenGate, and character-set encoding of JSON data. It covers the use of SQL/JSON condition `is json` as a check constraint to ensure that the data in a column is well-formed JSON data.

- [Overview of Storing and Managing JSON Data](#)
This overview describes: (1) data types for JSON columns, (2) LOB storage considerations for JSON data, and (3) ensuring that JSON columns contain well-formed JSON data.
- [Creating a Table With a JSON Column](#)
You can create a table that has JSON columns. You use SQL condition `is json` as a check constraint to ensure that data inserted into a column is (well-formed) JSON data. Oracle recommends that you *always* use an `is_json` check constraint when you create a column intended for JSON data.
- [SQL/JSON Conditions IS JSON and IS NOT JSON](#)
SQL/JSON conditions `is json` and `is not json` are complementary. They test whether their argument is syntactically correct, that is, *well-formed*, JSON data. You can use them in a `CASE` expression or the `WHERE` clause of a `SELECT` statement. You can use `is json` in a check constraint.
- [Character Sets and Character Encoding for JSON Data](#)
JSON data always uses the Unicode character set. In this respect, JSON data is simpler to use than XML data. This is an important part of the JSON Data Interchange Format (RFC 4627). For JSON data processed by Oracle Database, any needed character-set conversions are performed automatically.
- [Partitioning JSON Data](#)
You can partition a table using a JSON virtual column as the partitioning key. The virtual column is extracted from a JSON column using SQL/JSON function `json_value`.
- [Replication of JSON Data](#)
You can use Oracle GoldenGate to replicate tables that have columns containing JSON data.

3

Overview of Storing and Managing JSON Data

This overview describes: (1) data types for JSON columns, (2) LOB storage considerations for JSON data, and (3) ensuring that JSON columns contain well-formed JSON data.

Data Types for JSON Columns

You can store JSON data in Oracle Database using columns whose data types are `VARCHAR2`, `CLOB`, or `BLOB`. Whichever type you use, you can manipulate JSON data as you would manipulate any other data of those types. Storing JSON data using standard data types allows all features of Oracle Database, such as advanced replication, to work with tables containing JSON documents. The choice of which to use is typically motivated by the size of the JSON documents you need to manage:

- Use `VARCHAR2(4000)` if you are sure that your largest JSON documents do not exceed 4000 bytes (or characters)¹.

If you use Oracle Exadata then choosing `VARCHAR2(4000)` can improve performance by allowing the execution of some JSON operations to be pushed down to Exadata storage cells, for improved performance.

- Use `VARCHAR2(32767)` if you know that some of your JSON documents are larger than 4000 bytes (or characters) and you are sure that none of the documents exceeds 32767 bytes (or characters)¹.

With `VARCHAR2(32767)`, the first roughly 3.5K bytes (or characters) of a document is *stored in line*, as part of the table row. This means that the added cost of using `VARCHAR2(32767)` instead of `VARCHAR2(4000)` applies only to those documents that are larger than about 3.5K. If most of your documents are smaller than this then you will likely notice little performance difference from using `VARCHAR2(4000)`.

If you use Oracle Exadata then push-down is enabled for any documents that are stored in line.

- Use `BLOB` (binary large object) or `CLOB` (character large object) storage if you know that you have some JSON documents that are larger than 32767 bytes (or characters)¹.

Ensure That JSON Columns Contain Well-Formed JSON Data

You can use SQL/JSON condition `is json` to check whether or not some JSON data is well formed. Oracle strongly recommends that you apply an `is json` check constraint to any JSON column, unless you expect some rows to contain something other than well-formed JSON data.

¹ Whether the limit is expressed in bytes or characters is determined by session parameter `NLS_LENGTH_SEMANTICS`.

The overhead of parsing JSON is such that evaluating the condition should not have a significant impact on insert and update performance, and omitting the constraint means you cannot use the simple dot-notation syntax to query the JSON data.

What constitutes well-formed JSON data is a gray area. In practice, it is common for JSON data to have some characteristics that do not strictly follow the standard definition. You can control which syntax you require a given column of JSON data to conform to: the standard definition (strict syntax) or a JavaScript-like syntax found in common practice (lax syntax). The default SQL/JSON syntax for Oracle Database is *lax*. Which kind of syntax is used is controlled by condition `is json`. Applying an `is json` check constraint to a JSON column thus enables the use of lax JSON syntax, by default.

Considerations When Using LOB Storage for JSON Data

Oracle recommends that you use `BLOB`, not `CLOB` storage.

This is particularly relevant if the database character set is the Oracle-recommended value of `AL32UTF8`. In `AL32UTF8` databases `CLOB` instances are stored using the `UCS2` character set, which means that each character requires two bytes. This doubles the storage needed for a document if most of its content consists of characters that are represented using a single byte in character set `AL32UTF8`.

Even in cases where the database character set is not `AL32UTF8`, choosing `BLOB` over `CLOB` storage has the advantage that it avoids the need for character-set conversion when storing the JSON document (see [Character Sets and Character Encoding for JSON Data](#)).

When using large objects (LOBs), Oracle recommends that you do the following:

- Use the clause `LOB (COLUMN_NAME) STORE AS (CACHE)` in your `CREATE TABLE` statement, to ensure that read operations on the JSON documents are optimized using the database buffer cache.
- Use SecureFiles LOBs. Consider also using Oracle Advanced Compression, to reduce the storage space needed for your JSON data. If you use compression then Oracle recommends option Medium Compression, which provides a good balance between space savings and performance.

SQL/JSON functions and conditions work with JSON data without any special considerations, whether the data is stored as `BLOB` or `CLOB`. From an application-development perspective, the API calls for working with `BLOB` content are nearly identical to those for working with `CLOB` content.

A downside of choosing `BLOB` storage over `CLOB` (for JSON or any other kind of data) is that it is sometimes more difficult to work with `BLOB` content using command-line tools such as `SQL*Plus`. For instance:

- When selecting data from a `BLOB` column, if you want to view it as printable text then you must use SQL function `to_clob`.
- When performing insert or update operations on a `BLOB` column, you must explicitly convert character strings to `BLOB` format using SQL function `rawtohex`.²

² The return value of SQL function `rawtohex` is limited to 32767 bytes. The value is truncated to remove any converted data beyond this length.

 **See Also:**

- *Oracle Database SQL Language Reference* for information about SQL function `to_clob`
- *Oracle Database SQL Language Reference* for information about SQL function `rawtohex`

Considerations When Using a Client To Retrieve JSON Data As a LOB Instance

If you use a client, such as Oracle Call Interface (OCI) or Java Database Connectivity (JDBC), to retrieve JSON data from the database in LOB batches, consider using one of these optimizations:

- When using the *LOB data interface*:
 - In OCI, use data types `SQLT_BIN` and `SQLT_CHR`, for BLOB and CLOB data, respectively.
 - In JDBC, use methods `getBytes()` and `getString()`, for BLOB and CLOB data, respectively.
- When using the *LOB locator interface*:
 - Set the LOB prefetch size to a large value, such as 65 KB.
 - Free the fetched LOB after each row is read. Otherwise, performance can be reduced, and memory can leak.

This example illustrates how to use the LOB locator interface with JDBC.

```
Statement stmt = conn.createStatement(); // conn is a java.sql.Connection
stmt.setFetchSize(1000); // Set batch fetch size to 1000 rows
// set LOB prefetch size to be 65 KB
((OraclePreparedStatement)stmt).setLobPrefetchSize(65000);
// Query the JSON data in column jdoc of table myTab, returning JSON as a BLOB
instance
rset = stmt.executeQuery(
    "SELECT id, json_serialize(jdoc RETURNING BLOB) FROM myTab");
while (rset.next()) { // Iterate over the returned rows
    Blob blob = rset.getBlob (2);
    // Do something with the BLOB instance for the row
    blob.free(); // Free the LOB at the end of each iteration
}
rset.close();
```

Related Topics

- [Character Sets and Character Encoding for JSON Data](#)
JSON data always uses the Unicode character set. In this respect, JSON data is simpler to use than XML data. This is an important part of the JSON Data Interchange Format (RFC 4627). For JSON data processed by Oracle Database, any needed character-set conversions are performed automatically.

- [Overview of Inserting, Updating, and Loading JSON Data](#)
You can use standard database APIs to insert or update JSON data in Oracle Database. You can work directly with JSON data contained in file-system files by creating an external table that exposes it to the database. You can use JSON Merge Patch to update a JSON document.
- [Simple Dot-Notation Access to JSON Data](#)
Dot notation is designed for easy, general use and common use cases. Queries of JSON data that use dot-notation syntax return JSON values whenever possible.

**See Also:**

Oracle Database SecureFiles and Large Objects Developer's Guide

4

Creating a Table With a JSON Column

You can create a table that has JSON columns. You use SQL condition `is json` as a check constraint to ensure that data inserted into a column is (well-formed) JSON data. Oracle recommends that you *always* use an `is_json` check constraint when you create a column intended for JSON data.

[Example 4-1](#) and [Example 4-2](#) illustrate this. They create and fill a table that holds data used in examples elsewhere in this documentation.

For brevity, only two rows of data (one JSON document) are inserted in [Example 4-2](#).

Note:

SQL/JSON conditions `IS JSON` and `IS NOT JSON` return true or false for any non-NULL SQL value. But they both return unknown (neither true nor false) for SQL NULL. When used in a check constraint, they do not prevent a SQL NULL value from being inserted into the column. (But when used in a SQL WHERE clause, SQL NULL is never returned.)

It is true that a check constraint can reduce performance for data insertion. If you are sure that your application inserts only well-formed JSON data into a particular column, then consider *disabling* the check constraint, but *do not drop* the constraint.

See Also:

- [Loading External JSON Data](#) for the creation of the full table `j_purchaseorder`
- *Oracle Database SQL Language Reference* for information about `CREATE TABLE`

Example 4-1 Using IS JSON in a Check Constraint to Ensure JSON Data is Well-Formed

```
CREATE TABLE j_purchaseorder
(id          VARCHAR2 (32) NOT NULL PRIMARY KEY,
date_loaded TIMESTAMP (6) WITH TIME ZONE,
po_document VARCHAR2 (23767)
CONSTRAINT ensure_json CHECK (po_document IS JSON));
```

Example 4-2 Inserting JSON Data Into a VARCHAR2 JSON Column

```

INSERT INTO j_purchaseorder
VALUES (
  SYS_GUID(),
  to_date('30-DEC-2014'),
  '{"PONumber"          : 1600,
   "Reference"         : "ABULL-20140421",
   "Requestor"        : "Alexis Bull",
   "User"              : "ABULL",
   "CostCenter"       : "A50",
   "ShippingInstructions" :
     {"name"          : "Alexis Bull",
      "Address"       : {"street"   : "200 Sporting Green",
                        "city"     : "South San Francisco",
                        "state"    : "CA",
                        "zipCode"  : 99236,
                        "country"  : "United States of America"},
      "Phone"        : [{"type"    : "Office", "number" : "909-555-7307"},
                        {"type"    : "Mobile", "number" : "415-555-1234"}]},
   "Special Instructions" : null,
   "AllowPartialShipment" : true,
   [{"ItemNumber" : 1,
     "Part"       : {"Description" : "One Magic Christmas",
                    "UnitPrice"   : 19.95,
                    "UPCCode"     : 13131092899},
     "Quantity"   : 9.0},
    {"ItemNumber" : 2,
     "Part"       : {"Description" : "Lethal Weapon",
                    "UnitPrice"   : 19.95,
                    "UPCCode"     : 85391628927},
     "Quantity"   : 5.0}]]');

INSERT INTO j_purchaseorder
VALUES (
  SYS_GUID(),
  to_date('30-DEC-2014'),
  '{"PONumber"          : 672,
   "Reference"         : "SBELL-20141017",
   "Requestor"        : "Sarah Bell",
   "User"              : "SBELL",
   "CostCenter"       : "A50",
   "ShippingInstructions" : {"name"          : "Sarah Bell",
                            "Address"       : {"street"   : "200 Sporting Green",
                                                "city"     : "South San Francisco",
                                                "state"    : "CA",
                                                "zipCode"  : 99236,
                                                "country"  : "United States of America"},
                            "Phone"        : "983-555-6509"},
   "Special Instructions" : "Courier",
   "LineItems"          :
     [{"ItemNumber" : 1,
       "Part"       : {"Description" : "Making the Grade",
                      "UnitPrice"   : 20,
                      "UPCCode"     : 27616867759},
       "Quantity"   : 1}]]');

```

```

"Quantity" : 8.0},
{"ItemNumber" : 2,
 "Part" : {"Description" : "Nixon",
           "UnitPrice" : 19.95,
           "UPCCode" : 717951002396},
 "Quantity" : 5},
{"ItemNumber" : 3,
 "Part" : {"Description" : "Eric Clapton: Best Of 1981-1999",
           "UnitPrice" : 19.95,
           "UPCCode" : 75993851120},
 "Quantity" : 5.0}}}')';

```

- [Determining Whether a Column Necessarily Contains JSON Data](#)

How can you tell whether a given column of a table or view is well-formed JSON data? Whenever this is the case, the column is listed in the following static data dictionary views: `DBA_JSON_COLUMNS`, `USER_JSON_COLUMNS`, and `ALL_JSON_COLUMNS`.

4.1 Determining Whether a Column Necessarily Contains JSON Data

How can you tell whether a given column of a table or view is well-formed JSON data? Whenever this is the case, the column is listed in the following static data dictionary views: `DBA_JSON_COLUMNS`, `USER_JSON_COLUMNS`, and `ALL_JSON_COLUMNS`.

Each of these views lists the column name, data type, and format (`TEXT` or `BINARY`); the table or view name (column `TABLE_NAME`); and whether the object is a table or a view (column `OBJECT_TYPE`).

For a table column to be considered JSON data it must have an `is json` check constraint. But in the case of a view, any one of the following criteria suffices for a column to be considered JSON data:

- The underlying data has an `is json` check constraint.
- The column results from the use of SQL/JSON function `json_query`.
- The column results from the use of a JSON generation function, such as `json_object`.
- The column results from the use of SQL function `treat` with keywords `AS JSON`.

If an `is json` check constraint determines that a table column is JSON data, and if that constraint is later *deactivated*, the column remains listed in the views. If the check constraint is *dropped* then the column is removed from the views.

Note:

If a check constraint *combines* condition `is json` with another condition using logical condition `OR`, then the column is *not* listed in the views. In this case, it is not certain that data in the column is JSON data. For example, the constraint `jcol is json OR length(jcol) < 1000` does *not* ensure that the data in column `jcol` is JSON data.

 **See Also:**

Oracle Database Reference for information about `ALL_JSON_COLUMNS` and the related data-dictionary views

5

SQL/JSON Conditions IS JSON and IS NOT JSON

SQL/JSON conditions `is json` and `is not json` are complementary. They test whether their argument is syntactically correct, that is, *well-formed*, JSON data. You can use them in a `CASE` expression or the `WHERE` clause of a `SELECT` statement. You can use `is json` in a check constraint.

If the argument is syntactically correct then `is json` returns `true` and `is not json` returns `false`.

If an error occurs during parsing then the error is not raised, and the data is considered to *not* be well-formed: `is json` returns `false`; `is not json` returns `true`. If an error occurs other than during parsing then that error is raised.

Well-formed data means syntactically correct data. JSON data stored textually can be well-formed in two senses, referred to as strict and lax syntax. In addition, for textual JSON data you can specify whether a JSON object can have duplicate fields (keys).

Whenever textual JSON data is generated inside the database it satisfies condition `is json` with keyword `STRICT`. This includes generation in these ways:

- Using a SQL/JSON generation function (unless you specify keyword `STRICT` with either `FORMAT JSON` or `TREAT AS JSON`, which means that you declare that the data is JSON data; you vouch for it, so its well-formedness is not checked)
- Using SQL function `json_serialize`
- Using SQL function `to_clob`, `to_blob`, or `to_varchar2` on a PL/SQL DOM
- Using SQL/JSON function `json_query`
- Using SQL/JSON function `json_table` with `FORMAT JSON`
- [Unique Versus Duplicate Fields in JSON Objects](#)
By default, field names need not be unique for a given JSON object. But you can specify that particular JSON data is to be considered well-formed only if none of its objects have duplicate field names.
- [About Strict and Lax JSON Syntax](#)
The Oracle default syntax for JSON is lax. In particular: it reflects the JavaScript syntax for object fields; the Boolean and null values are not case-sensitive; and it is more permissive with respect to numerals, whitespace, and escaping of Unicode characters.
- [Specifying Strict or Lax JSON Syntax](#)
The default JSON syntax for Oracle Database is lax. Strict or lax syntax matters *only* for SQL/JSON conditions `is json` and `is not json`. All other SQL/JSON functions and conditions use lax syntax for interpreting input and strict syntax when returning output.

Related Topics

- [Creating a Table With a JSON Column](#)
You can create a table that has JSON columns. You use SQL condition `is json` as a check constraint to ensure that data inserted into a column is (well-formed) JSON data. Oracle recommends that you *always* use an `is_json` check constraint when you create a column intended for JSON data.

See Also:

Oracle Database SQL Language Reference for information about `is json` and `is not json`.

5.1 Unique Versus Duplicate Fields in JSON Objects

By default, field names need not be unique for a given JSON object. But you can specify that particular JSON data is to be considered well-formed only if none of its objects have duplicate field names.

The JSON standard does not specify whether field names must be unique for a given JSON object. This means that, a priori, a well-formed JSON object can have multiple members that have the same field name. This is the *default* behavior for handling JSON data in Oracle Database because checking for duplicate names takes additional time.

You can specify that particular JSON data is to be considered well-formed only if all of the objects it contains have unique field names, that is, no object has duplicate field names. You do this by using the keywords `WITH UNIQUE KEYS` with SQL/JSON condition `is json`.¹

If you do not specify `UNIQUE KEYS`, or if you use the keywords `WITHOUT UNIQUE KEYS`, then objects can have duplicate field names and still be considered well-formed.

The evaluation that Oracle Database employs always uses only one of the object members that have a given field name; any other members with the same field name are ignored. It is unspecified which of multiple such members is used.

Whether duplicate field names are allowed in well-formed JSON data is orthogonal to whether Oracle uses strict or lax syntax to determine well-formedness.

5.2 About Strict and Lax JSON Syntax

The Oracle default syntax for JSON is lax. In particular: it reflects the JavaScript syntax for object fields; the Boolean and null values are not case-sensitive; and it is more permissive with respect to numerals, whitespace, and escaping of Unicode characters.

Standard ECMA-404, the *JSON Data Interchange Format*, and ECMA-262, the *ECMAScript Language Specification*, define JSON syntax.

¹ An object field is sometimes called an object “key”.

According to these specifications, each JSON field and each string value must be enclosed in double quotation marks ("). Oracle supports this **strict JSON syntax**, but it is *not* the default syntax.

In JavaScript notation, a field used in an object literal can be, but need not be, enclosed in double quotation marks. It can also be enclosed in single quotation marks ('). Oracle also supports this **lax JSON syntax**, and it is the *default* syntax.

In addition, in practice, some JavaScript implementations (but not the JavaScript standard) allow one or more of the following:

- Case variations for keywords `true`, `false`, and `null` (for example, `TRUE`, `True`, `TrUe`, `fALSe`, `Null`).
- An extra comma (,) after the last element of an array or the last member of an object (for example, `[a, b, c,]`, `{a:b, c:d,}`).
- Numerals with one or more leading zeros (for example, `0042.3`).
- Fractional numerals that lack 0 before the decimal point (for example, `.14` instead of `0.14`).
- Numerals with no fractional part after the decimal point (for example, `342.` or `1.e27`).
- A plus sign (+) preceding a numeral, meaning that the number is non-negative (for example, `+1.3`).

This syntax too is allowed as part of the Oracle default (lax) JSON syntax. (See the JSON standard for the strict numeral syntax.)

In addition to the ASCII space character (U+0020), the JSON standard defines the following characters as insignificant (ignored) whitespace when used outside a quoted field or a string value:

- Tab, horizontal tab (HT, ^I, decimal 9, U+0009, \t)
- Line feed, newline (LF, ^J, decimal 10, U+000A, \n)
- Carriage return (CR, ^M, decimal 13, U+000D, \r)

The lax JSON syntax, however, treats *all* of the ASCII control characters (Control+0 through Control+31), as well as the ASCII space character (decimal 32, U+0020), as (insignificant) whitespace characters. The following are among the control characters:

- Null (NUL, ^@, decimal 0, U+0000, \0)
- Bell (BEL, ^G, decimal 7, U+0007, \a)
- Vertical tab (VT, ^K, decimal 11, U+000B)
- Escape (ESC, ^[, decimal 27, U+001B, \e)
- Delete (DEL, ^?, decimal 127, U+007F)

An ASCII space character (U+0020) is the only whitespace character allowed, unescaped, within a quoted field or a string value. This is true for both the lax and strict JSON syntaxes.

For both strict and lax JSON syntax, quoted object field and string values can contain any Unicode character, but some of them must be escaped, as follows:

- ASCII control characters are not allowed, except for those represented by the following escape sequences: `\b` (backspace), `\f` (form feed), `\n` (newline, line feed), `\r` (carriage return), and `\t` (tab, horizontal tab).
- Double quotation mark (`"`), slash (`/`), and backslash (`\`) characters must also be escaped (preceded by a backslash): `\"`, `\/`, and `\\`, respectively.

In the lax JSON syntax, an object field that is *not* quoted can contain any Unicode character except whitespace and the JSON structural characters — left and right brackets (`[`, `]`) and curly braces (`{`, `}`), colon (`:`), and comma (`,`), but escape sequences are not allowed.

Any Unicode character can also be included in a name or string by using the ASCII escape syntax `\u` followed by the four ASCII hexadecimal digits that represent the Unicode code point.

Note that other Unicode characters that are not printable or that might appear as whitespace, such as a no-break space character (U+00A0), are *not* considered whitespace for either the strict or the lax JSON syntax.

[Table 5-1](#) shows some examples of JSON syntax.

Table 5-1 JSON Object Field Syntax Examples

Example	Well-Formed?
<code>"part number": 1234</code>	Lax and strict: yes. Space characters are allowed.
<code>part number: 1234</code>	Lax (and strict): no . Whitespace characters, including space characters, are not allowed in unquoted names.
<code>"part\tnumber": 1234</code>	Lax and strict: yes. Escape sequence for tab character is allowed.
<code>"part number": 1234</code>	Lax and strict: no . Unescaped tab character is not allowed. Space is the only unescaped whitespace character allowed.
<code>"\"part\"number": 1234</code>	Lax and strict: yes. Escaped double quotation marks are allowed, if name is quoted.
<code>\\"part\"number: 1234</code>	Lax and strict: no . Name must be quoted.
<code>'\"part\"number': 1234</code>	Lax: yes, strict: no . Single-quoted names (object fields and strings) are allowed for lax syntax only. Escaped double quotation mark is allowed in a quoted name.
<code>"pärt : number":1234</code>	Lax and strict: yes. Any Unicode character is allowed in a quoted name. This includes whitespace characters and characters, such as colon (<code>:</code>), that are structural in JSON.
<code>part:number:1234</code>	Lax (and strict): no . Structural characters are not allowed in unquoted names.

Related Topics

- [JSON Syntax and the Data It Represents](#)
JSON (and JavaScript) values, scalars, objects, and arrays are described.

 **See Also:**

- IETF RFC4627 and ECMA 404 for the syntax of JSON Data Interchange Format
- ECMA International and JSON.org for more information about JSON and JavaScript

5.3 Specifying Strict or Lax JSON Syntax

The default JSON syntax for Oracle Database is lax. Strict or lax syntax matters *only* for SQL/JSON conditions `is json` and `is not json`. All other SQL/JSON functions and conditions use lax syntax for interpreting input and strict syntax when returning output.

If you need to be sure that particular JSON input data has strictly correct syntax, then check it first using `is json` or `is not json`.

You specify that data is to be checked as strictly well-formed according to the JSON standard by appending (`STRICT`) (parentheses included) to an `is json` or an `is not json` expression.

[Example 5-1](#) illustrates this. It is identical to [Example 4-1](#) except that it uses (`STRICT`) to ensure that all data inserted into the column is well-formed according to the JSON standard.

 **See Also:**

Oracle Database SQL Language Reference for information about `CREATE TABLE`

Example 5-1 Using IS JSON in a Check Constraint to Ensure JSON Data is Strictly Well-Formed (Standard)

```
CREATE TABLE j_purchaseorder
(id          VARCHAR2 (32) NOT NULL PRIMARY KEY,
 date_loaded  TIMESTAMP (6) WITH TIME ZONE,
 po_document VARCHAR2 (32767)
 CONSTRAINT ensure_json CHECK (po_document IS JSON (STRICT)));
```

Related Topics

- [About Strict and Lax JSON Syntax](#)
The Oracle default syntax for JSON is lax. In particular: it reflects the JavaScript syntax for object fields; the Boolean and null values are not case-sensitive; and it is more permissive with respect to numerals, whitespace, and escaping of Unicode characters.

6

Character Sets and Character Encoding for JSON Data

JSON data always uses the Unicode character set. In this respect, JSON data is simpler to use than XML data. This is an important part of the JSON Data Interchange Format (RFC 4627). For JSON data processed by Oracle Database, any needed character-set conversions are performed automatically.

Oracle Database uses UTF-8 internally when it processes JSON data (parsing, querying). If the data that is input to such processing, or the data that is output from it, must be in a different character set from UTF-8, then character-set conversion is carried out accordingly.

Character-set conversion can affect performance. And in some cases it can be lossy. Conversion of input data to UTF-8 is a lossless operation, but conversion to output can result in *information loss* in the case of characters that cannot be represented in the output character set.

If your JSON data is stored in the database as Unicode then no character-set conversion is needed for storage or retrieval. This is the case if the database character set is AL32UTF8 (Unicode UTF-8) or if your JSON data is stored in CLOB instances that have character set AL16UTF16. Oracle recommends that you *use AL32UTF8 as the database character set* if at all possible.

Regardless of the database character set, JSON data that is stored using data type BLOB never undergoes character-set conversion for storage or retrieval. JSON data can be stored using data type BLOB as AL32UTF8, AL16UTF16, or AL16UTF16LE.

If you *transform* JSON data using SQL/JSON functions or PL/SQL methods and you return the result of the transformation using data type BLOB then the result is encoded as AL32UTF8. This is true even if the input BLOB data uses another Unicode encoding.

For example, if you use function `json_query` to extract some JSON data from BLOB input and return the result as BLOB, it is returned using AL32UTF8.

Lossy character-set conversion can occur if application of a SQL/JSON function or a PL/SQL method specifies a return data type of VARCHAR2 or CLOB and the database character set is not Unicode. If input JSON data was stored in a BLOB instance then even if it is ultimately written again as BLOB, if some of it was temporarily changed to VARCHAR2 or CLOB then the resulting BLOB data can suffer from lossy conversion.

Related Topics

- [Overview of Storing and Managing JSON Data](#)

This overview describes: (1) data types for JSON columns, (2) LOB storage considerations for JSON data, and (3) ensuring that JSON columns contain well-formed JSON data.

 **See Also:**

- Unicode.org for information about Unicode
- IETF RFC4627 and ECMA 404 for the JSON Data Interchange Format
- *Oracle Database Migration Assistant for Unicode Guide* for information about using different character sets with the database
- *Oracle Database Globalization Support Guide* for information about character-set conversion in the database

7

Partitioning JSON Data

You can partition a table using a JSON virtual column as the partitioning key. The virtual column is extracted from a JSON column using SQL/JSON function `json_value`.

Partition on a Non-JSON Column When Possible

You can partition a table using a JSON virtual column, but it is generally preferable to use a non-JSON column. A partitioning key specifies which partition a new row is inserted into. A partitioning key defined as a JSON virtual column uses SQL/JSON function `json_value`, and the partition-defining `json_value` expression is *executed each time a row is inserted*. This can be costly, especially for insertion of large JSON documents.

Rules for Partitioning a Table Using a JSON Virtual Column

- The virtual column that serves as the partitioning key must be defined using SQL/JSON function `json_value`.
- The data type of the virtual column is defined by the `RETURNING` clause used for the `json_value` expression.
- The `json_value` path used to extract the data for the virtual column must not contain any predicates. (The path must be streamable.)
- The JSON column referenced by the expression that defines the virtual column can have an `is json` check constraint, but it *need not* have such a constraint.



See Also:

Oracle Database SQL Language Reference for information about `CREATE TABLE`

Example 7-1 Creating a Partitioned Table Using a JSON Virtual Column

This example creates table `j_purchaseorder_partitioned`, which is partitioned using virtual column `po_num_vc`. That virtual column references JSON column `po_document` (which uses CLOB storage). The `json_value` expression that defines the virtual column extracts JSON field `PONumber` from `po_document` as a number. Column `po_document` does *not* have an `is json` check constraint.

```
CREATE TABLE j_purchaseorder_partitioned
(id VARCHAR2 (32) NOT NULL PRIMARY KEY,
 date_loaded TIMESTAMP (6) WITH TIME ZONE,
 po_document CLOB,
 po_num_vc NUMBER GENERATED ALWAYS AS
  (json_value (po_document, '$.PONumber' RETURNING NUMBER)))
LOB (po_document) STORE AS (CACHE)
```

```
PARTITION BY RANGE (po_num_vc)
(PARTITION p1 VALUES LESS THAN (1000),
 PARTITION p2 VALUES LESS THAN (2000));
```


8

Replication of JSON Data

You can use Oracle GoldenGate to replicate tables that have columns containing JSON data.

Be aware that Oracle GoldenGate requires tables that are to be replicated to have a nonvirtual primary key column; the *primary key column cannot be virtual*.

All *indexes* on the JSON data will be replicated also. However, you must execute, on the replica database, any Oracle Text operations that you use to maintain a JSON search index. Here are examples of such procedures:

- `CTX_DDL.sync_index`
- `CTX_DDL.optimize_index`

See Also:

- *Oracle GoldenGate* for information about Oracle GoldenGate
- *Oracle Text Reference* for information about `CTX_DDL.sync_index`
- *Oracle Text Reference* for information about `CTX_DDL.optimize_index`

Part III

Insert, Update, and Load JSON Data

The usual ways to insert, update, and load data in Oracle Database work with JSON data. You can also create an external table from the content of a JSON dump file.

- [Overview of Inserting, Updating, and Loading JSON Data](#)
You can use standard database APIs to insert or update JSON data in Oracle Database. You can work directly with JSON data contained in file-system files by creating an external table that exposes it to the database. You can use JSON Merge Patch to update a JSON document.
- [Loading External JSON Data](#)
You can create a database table of JSON data from the content of a JSON dump file.
- [Updating a JSON Document with JSON Merge Patch](#)
You can use Oracle SQL function `json_mergepatch` to update specific portions of a JSON document. You pass it a JSON Merge Patch document, which specifies the changes to make to a specified JSON document. JSON Merge Patch is an IETF standard.

9

Overview of Inserting, Updating, and Loading JSON Data

You can use standard database APIs to insert or update JSON data in Oracle Database. You can work directly with JSON data contained in file-system files by creating an external table that exposes it to the database. You can use JSON Merge Patch to update a JSON document.

Use Standard Database APIs to Insert or Update JSON Data

Because JSON data is stored using standard SQL data types, all of the standard database APIs used to insert or update `VARCHAR2` and large-object (LOB) columns can be used for columns containing JSON documents. To these APIs, a stored JSON document is nothing more than a string of characters.

You specify that a JSON column must contain only well-formed JSON data by using SQL condition `is json` as a check constraint. The database handles this check constraint the same as any other check constraint — it enforces rules about the content of the column. Working with a column of type `VARCHAR2`, `BLOB`, or `CLOB` that contains JSON documents is thus no different from working with any other column of that type.

Update operations on a document in a JSON column require the replacement of the entire document. You can make fine-grained modifications to a JSON document, but when you need to save the changes to disk the entire updated document is written.

Inserting a JSON document into a JSON column is straightforward if the column is of data type `VARCHAR2` or `CLOB` — see [Example 4-2](#). The same is true of updating such a column.

But if you use a command-line tool such as SQL*Plus to insert data into a JSON column of type `BLOB`, or to update such data, then *you must convert the JSON data properly to binary format*. [Example 9-1](#) is a partial example of this. It assumes that table `my_table` has a JSON column, `json_doc`, which uses `BLOB` storage.

Example 9-1 Inserting JSON Data Into a BLOB Column

The textual JSON data being inserted (shown as partially elided literal data, `{...}`) contains characters in the database character set, which is `WE8MSWIN1252`. The data is passed to PL/SQL function `UTL_RAW.cast_to_raw`, which casts the data type to `RAW`. That result is then passed to function `UTL_RAW.convert`, which converts it to character set `AL32UTF8`.

```
INSERT INTO my_table (json_doc)
VALUES (UTL_RAW.convert(UTL_RAW.cast_to_raw('{...}'),
                        'AL32UTF8',
                        'WE8MSWIN1252'));
```

Use an External Table to Work With JSON Data in File-System Files

External tables make it easy to access JSON documents that are stored as separate files in a file system. Each file can be exposed to Oracle Database as a row in an external table. An external table can also provide access to the content of a dump file produced by a NoSQL database. You can use an external table of JSON documents to, in effect, *query the data in file-system files directly*. This can be useful if you need only process the data from all of the files in a one-time operation.

But if you instead need to make multiple queries of the documents, and especially if different queries select data from different rows of the external table (different documents), then for better performance consider copying the data from the external table into an ordinary database table, using an `INSERT AS SELECT` statement — see [Example 10-4](#). Once the JSON data has been loaded into a JSON column of an ordinary table, you can index the content, and then you can efficiently query the data in a repetitive, selective way.

Use JSON Merge Patch To Update a JSON Document

You can use Oracle SQL function `json_mergepatch` or PL/SQL object-type method `json_mergepatch()` to update specific portions of a JSON document. In both cases you provide a JSON Merge Patch document, which declaratively specifies the changes to make to a specified JSON document. JSON Merge Patch is an IETF standard.

See Also:

- IETF RFC7396 for the definition of JSON Merge Patch
- *Oracle Database SQL Language Reference* for information about SQL function `json_mergepatch`

Note:

In addition to the usual ways to insert, update, and load data, you can use *Simple Oracle Document Access* (SODA) APIs. SODA is designed for schemaless application development without knowledge of relational database features or languages such as SQL and PL/SQL. It lets you create and store collections of documents of any kind (not just JSON), retrieve them, and query them, without needing to know how the documents are stored in the database. SODA also provides query features that are specific for JSON documents.

There are two implementations of SODA:

- SODA for Java — Java classes that represent database, collection, and document.
- SODA for REST — SODA operations as representational state transfer (REST) requests, using any language capable of making HTTP calls.

For information about SODA see [Oracle as a Document Store](#).

Related Topics

- [Loading External JSON Data](#)
You can create a database table of JSON data from the content of a JSON dump file.
- [Updating a JSON Document with JSON Merge Patch](#)
You can use Oracle SQL function `json_mergepatch` to update specific portions of a JSON document. You pass it a JSON Merge Patch document, which specifies the changes to make to a specified JSON document. JSON Merge Patch is an IETF standard.
- [Creating a Table With a JSON Column](#)
You can create a table that has JSON columns. You use SQL condition `is json` as a check constraint to ensure that data inserted into a column is (well-formed) JSON data. Oracle recommends that you *always* use an `is_json` check constraint when you create a column intended for JSON data.
- [Overview of Storing and Managing JSON Data](#)
This overview describes: (1) data types for JSON columns, (2) LOB storage considerations for JSON data, and (3) ensuring that JSON columns contain well-formed JSON data.

See Also:

- [PL/SQL Object Types for JSON](#) for information about updating JSON data using PL/SQL object types
- *Oracle Database SQL Language Reference* for information about SQL function `rawtohex`

10

Loading External JSON Data

You can create a database table of JSON data from the content of a JSON dump file.

This topic shows how you can load a full table of JSON documents from the data in a JSON dump file, `$ORACLE_HOME/demo/schema/order_entry/PurchaseOrders.dmp`. The format of this file is compatible with the export format produced by common NoSQL databases, including Oracle NoSQL Database. Each row of the file contains a single JSON document represented as a JSON object.

You can query such an external table directly or, for better performance if you have multiple queries that target different rows, you can load an ordinary database table from the data in the external table.

[Example 10-1](#) creates a *database directory* that corresponds to file-system directory `$ORACLE_HOME/demo/schema/order_entry`. [Example 10-2](#) then uses this database directory to create and fill an *external table*, `json_dump_file_contents`, with the data from the dump file, `PurchaseOrders.dmp`. It bulk-fills the external table completely, copying all of the JSON documents to column `json_document`.

[Example 10-4](#) then uses an `INSERT AS SELECT` statement to copy the JSON documents from the external table to JSON column `po_document` of ordinary database table `j_purchaseorder`.

Because we chose `BLOB` storage for JSON column `json_document` of the external table, column `po_document` of the ordinary table must also be of type `BLOB`.

[Example 10-3](#) creates table `j_purchaseorder` with `BLOB` column `po_document`.

Note:

You need system privilege `CREATE ANY DIRECTORY` to create a database directory.

See Also:

- *Oracle Database Concepts* for overview information about external tables
- *Oracle Database Utilities* and *Oracle Database Administrator's Guide* for detailed information about external tables
- *Oracle Database Data Warehousing Guide*
- *Oracle Database SQL Language Reference* for information about `CREATE TABLE`

Example 10-1 Creating a Database Directory Object for Purchase Orders

You must replace `$ORACLE_HOME` here by its value.

```
CREATE OR REPLACE DIRECTORY order_entry_dir
AS '$ORACLE_HOME/demo/schema/order_entry';
```

Example 10-2 Creating an External Table and Filling It From a JSON Dump File

```
CREATE TABLE json_dump_file_contents (json_document BLOB)
  ORGANIZATION EXTERNAL (TYPE ORACLE_LOADER DEFAULT DIRECTORY
order_entry_dir
                        ACCESS PARAMETERS
                        (RECORDS DELIMITED BY 0x'0A'
                        DISABLE_DIRECTORY_LINK_CHECK
                        FIELDS (json_document CHAR(5000)))
                        LOCATION
(order_entry_dir:'PurchaseOrders.dmp'))
  PARALLEL
  REJECT LIMIT UNLIMITED;
```

Example 10-3 Creating a Table With a BLOB JSON Column

Table `j_purchaseorder` has primary key `id` and JSON column `po_document`, which is stored using data type `BLOB`. The *LOB cache option* is turned on for that column.

```
DROP TABLE j_purchaseorder;

CREATE TABLE j_purchaseorder
(id          VARCHAR2 (32) NOT NULL PRIMARY KEY,
 date_loaded  TIMESTAMP (6) WITH TIME ZONE,
 po_document  BLOB
  CONSTRAINT ensure_json CHECK (po_document IS JSON))
LOB (po_document) STORE AS (CACHE);
```

Example 10-4 Copying JSON Data From an External Table To a Database Table

```
INSERT INTO j_purchaseorder (id, date_loaded, po_document)
  SELECT SYS_GUID(), SYSTIMESTAMP, json_document FROM
json_dump_file_contents
  WHERE json_document IS JSON;
```

11

Updating a JSON Document with JSON Merge Patch

You can use Oracle SQL function `json_mergepatch` to update specific portions of a JSON document. You pass it a JSON Merge Patch document, which specifies the changes to make to a specified JSON document. JSON Merge Patch is an IETF standard.

You can use `json_mergepatch` in an `UPDATE` statement, to update the documents in a JSON column. [Example 11-3](#) illustrates this.

You can use `json_mergepatch` in a `SELECT` list, to modify the selected documents. The modified documents can be returned or processed further. [Example 11-4](#) illustrates this.

JSON Merge Patch is suitable for updating JSON documents that primarily use *objects* for their structure and do not make use of explicit `null` values. You cannot use it to add, remove, or change array elements (except by explicitly replacing the whole array). And you cannot use it to set the value of a field to `null`.

JSON Merge Patch acts a bit like a UNIX `patch` utility: you give it (1) a *source* document to patch and (2) a *patch* document that specifies the changes to make, and it returns a copy of the source document updated (patched). The patch document specifies the differences between the source and the result documents. For UNIX `patch` the differences are in the form of UNIX `diff` utility output. For JSON Merge Patch both source and patch are JSON documents.

You can think of JSON Merge Patch as *merging* the contents of the source and the patch. When merging two objects, one from source and one from patch, a member with a field that is in one object but not in the other is kept in the result. An exception is that a patch member with field value is `null` is ignored when the source object has no such field.

When merging object members that have the same field:

- If the patch field value is `null` then the field is dropped from the source — it is not included in the result.
- Otherwise, the field is kept in the result, but its value is the *result of merging* the source field value with the patch field value. That is, the merging operation in this case is recursive — it dives down into fields whose values are themselves objects.

A little more precisely, JSON Merge Patch acts as follows:

- If the *patch* is *not* a JSON object then *replace* the source by the patch.
- Otherwise (the patch is an object), do the following:
 1. If the *source* is *not* an object then act as if it were the empty object (`{}`).
 2. Iterate over the (*p-field:p-value*) members of the patch object.
 - If the *p-value* of the patch member is `null` then *remove* the corresponding member from the source.

- Otherwise, **recurse**: *Replace* the value of the corresponding source field with the *result of merge-patching* that value (as the next source) with the *p-value* (as the next patch).

If a patch field value of `null` did not have a special meaning (remove the corresponding source member with that field) then you could use it as a field value to set the corresponding source field value to `null`. The special removal behavior means you *cannot* set a source field value to `null`.

Examples:

- Patch member `"PONumber":99999` overrides a source member with field `PONumber`, *replacing its value* with the patch-specified value, `99999`.

```
json_mergepatch('{ "User": "ABULL", "PONumber": 1600 }',
  '{ "PONumber": 99999 }') results in { "User": "ABULL", "PONumber": 99999 }.
```

- Patch member `"tracking":123456` overrides a missing source member with field `tracking`, *adding* that patch member to the result. And source member `"PONumber":1600` overrides a missing patch member with field `PONumber` — it is kept in the result.

```
json_mergepatch('{ "PONumber": 1600 }', '{ "tracking": 123456 }') results in
{ "PONumber": 1600, "tracking": 123456 }.
```

- Patch member `"Reference":null` overrides a source member with field `Reference`, *removing* it from the result.

```
json_mergepatch('{ "PONumber": 1600, "Reference": "ABULL-20140421" }',
  '{ "Reference": null }') results in { "PONumber": 1600 }.
```

- Patch value `[1,2,3]` overrides the corresponding source value, `[4,5,6]`, *replacing* it.

```
json_mergepatch('{ "PONumber": 1600, "LineItems": [1, 2, 3] }',
  '{ "LineItems": [4, 5, 6] }') results in { "PONumber": 1600, "LineItems": [4, 5, 6] }.
```

Note:

The merge-patch procedure — in particular the fact that there is no recursive behavior for a non-object patch — means that you *cannot* add, remove, or replace values of an array individually. To make such a change you must *replace the whole array*. For example, if the source document has a member `Phone: ["999-555-1212", "415-555-1234"]` then to remove the second phone number you can use a patch whose content has a member `"Phone": ["999-555-1212"]`.

See Also:

- IETF RFC7396 for the definition of JSON Merge Patch
- *Oracle Database SQL Language Reference* for information about SQL function `json_mergepatch`

Example 11-1 A JSON Merge Patch Document

If applied to the document shown in [Example 2-1](#), this JSON Merge Patch document does the following:

- Adds member "Category" : "Platinum".
- Removes the member with field ShippingInstructions.
- Replaces the value of field Special Instructions with the string "Contact User SBELL".
- Replaces the value of field LineItems with the empty array, []
- Replaces member "AllowPartialShipment" : null with member "Allow Partial Shipment" : false (in effect *renaming* the field, since the field value was already false).

```
{ "Category" : "Platinum",
  "ShippingInstructions" : null,
  "Special Instructions" : "Contact User SBELL",
  "LineItems" : [],
  "AllowPartialShipment" : null,
  "Allow Partial Shipment" : false }
```

Example 11-2 A Merge-Patched JSON Document

This example shows the document that results from merge-patching the document in [Example 2-1](#) with the patch of [Example 11-1](#).

```
{ "PONumber" : 1600,
  "Reference" : "ABULL-20140421",
  "Requestor" : "Alexis Bull",
  "User" : "ABULL",
  "CostCenter" : "A50",
  "Special Instructions" : "Contact User SBELL",
  "Allow Partial Shipment" : false,
  "LineItems" : [],
  "Category" : "Platinum" }
```

Example 11-3 Updating a JSON Column Using JSON Merge Patch

This example updates all documents in `j_purchaseorder.po_document`, removing field "Special Instructions".

```
UPDATE j_purchaseorder SET po_document =
  json_mergepatch(po_document, '{"Special Instructions':null}');
```

Example 11-4 Updating Selected JSON Documents On the Fly

This example selects all documents in `j_purchaseorder.po_document`, returning pretty-printed, updated copies of them, where field "Special Instructions" has been removed.

```
SELECT json_mergepatch(po_document, '{"Special Instructions":null}'
                     RETURNING CLOB PRETTY)
FROM j_purchaseorder;
```

Related Topics

- [Overview of Inserting, Updating, and Loading JSON Data](#)
You can use standard database APIs to insert or update JSON data in Oracle Database. You can work directly with JSON data contained in file-system files by creating an external table that exposes it to the database. You can use JSON Merge Patch to update a JSON document.
- [Using PL/SQL Object Types for JSON](#)
Some examples of using PL/SQL object types for JSON are presented.
- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no `RETURNING` clause) are described here.

Part IV

Query JSON Data

You can query JSON data using a simple dot notation or, for more functionality, using SQL/JSON functions and conditions. You can create and query a *data guide* that summarizes the structure and type information of a set of JSON documents.

Because JSON data is stored in the database using standard data types (`VARCHAR2`, `BLOB`, and `CLOB`), SQL queries work with JSON data the same as with any other database data.

To query particular JSON fields, or to map particular JSON fields to SQL columns, you can use the SQL/JSON *path language*. In its simplest form a path expression consists of one or more field names separated by periods (`.`). More complex path expressions can contain filters and array indexes.

Oracle provides two ways of querying JSON content:

- *A dot-notation syntax*, which is essentially a table alias, followed by a JSON column name, followed by one or more field names — all separated by periods (`.`). An array step can follow each of the field names. This syntax is designed to be simple to use and to return JSON values whenever possible.
- *SQL/JSON functions and conditions*, which completely support the path language and provide more power and flexibility than is available using the dot-notation syntax. You can use them to create, query, and operate on JSON data stored in Oracle Database.
 - Condition `json_exists` tests for the existence of a particular value within some JSON data.
 - Conditions `is json` and `is not json` test whether some data is well-formed JSON data. The former is used especially as a check constraint.
 - Function `json_value` selects a scalar value from some JSON data, as a SQL value.
 - Function `json_query` selects one or more values from some JSON data, as a SQL string representing the JSON values. It is used especially to retrieve fragments of a JSON document, typically a JSON object or array.
 - Function `json_table` projects some JSON data as a virtual table, which you can also think of as an inline view.

Because the path language is part of the query language, no fixed schema is imposed on the data. This design supports *schemaless development*. A “schema”, in effect, gets defined on the fly at *query time*, by your specifying a given path. This is in contrast to the more usual approach with SQL of defining a schema (a set of table rows and columns) for the data at *storage time*.

Oracle SQL condition `json_equal` does not accept a path-expression argument. It just compares two JSON values and returns true if they are equal, false otherwise. For this comparison, insignificant whitespace and insignificant object member order are ignored. For example, JSON objects are equal if they have the same members,

regardless of their order. However, if either of two compared objects has one or more duplicate fields then the value returned by `json_equal` is unspecified.

You can generate and query a JSON *data guide*, to help you develop expressions for navigating your JSON content. A data guide can give you a deep understanding of the structure and type information of your JSON documents. Data guide information can be updated automatically, to keep track of new documents that you add.

- [Simple Dot-Notation Access to JSON Data](#)
Dot notation is designed for easy, general use and common use cases. Queries of JSON data that use dot-notation syntax return JSON values whenever possible.
- [SQL/JSON Path Expressions](#)
Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.
- [Clauses Used in SQL Query Functions and Conditions](#)
Clauses `RETURNING`, `wrapper`, `error`, and `empty-field` are described. Each is used in one or more of the SQL functions and conditions `json_value`, `json_query`, `json_table`, `json_serialize`, `json_mergepatch`, `is json`, `is not json`, `json_exists`, and `json_equal`.
- [SQL/JSON Condition JSON_EXISTS](#)
SQL/JSON condition `json_exists` lets you use a SQL/JSON path expression as a row filter, to select rows based on the content of JSON documents. You can use condition `json_exists` in a `CASE` expression or the `WHERE` clause of a `SELECT` statement.
- [SQL/JSON Function JSON_VALUE](#)
SQL/JSON function `json_value` selects JSON data and returns a SQL scalar or an instance of a user-defined SQL object type or SQL collection type (`varray`, nested table).
- [SQL/JSON Function JSON_QUERY](#)
SQL/JSON function `json_query` selects one or more values from JSON data and returns a string (`VARCHAR2`, `CLOB`, or `BLOB` instance) that represents the JSON values. You can thus use `json_query` to retrieve *fragments* of a JSON document.
- [SQL/JSON Function JSON_TABLE](#)
SQL/JSON function `json_table` projects specific JSON data to columns of various SQL data types. You use it to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.
- [Oracle SQL Function JSON_SERIALIZE](#)
Oracle SQL function `json_serialize` takes JSON data (of any SQL data type, `VARCHAR2`, `CLOB`, or `BLOB`) as input and returns a textual representation of it (as `VARCHAR2`, `CLOB`, or `BLOB` data). `VARCHAR2(4000)` is the default return type.
- [JSON Data Guide](#)
A JSON data guide lets you discover information about the structure and content of JSON documents stored in Oracle Database.



See Also:

Oracle Database SQL Language Reference for complete information about the syntax and semantics of the SQL/JSON functions and conditions

12

Simple Dot-Notation Access to JSON Data

Dot notation is designed for easy, general use and common use cases. Queries of JSON data that use dot-notation syntax return JSON values whenever possible.

The return value for a dot-notation query is *always a string* (data type `VARCHAR2(4000)`) representing JSON data. The content of the string depends on the targeted JSON data, as follows:

- If a *single* JSON value is targeted, then that value is the string content, whether it is a JSON scalar, object, or array.
- If *multiple* JSON values are targeted, then the string content is a JSON array whose elements are those values.

This behavior contrasts with that of SQL/JSON functions `json_value` and `json_query`, which you can use for more complex queries. They can return `NULL` or raise an error if the path expression you provide them does not match the queried JSON data. They accept optional clauses to specify the data type of the return value (`RETURNING` clause), whether or not to wrap multiple values as an array (`wrapper` clause), how to handle errors generally (`ON ERROR` clause), and how to handle missing JSON fields (`ON EMPTY` clause).

In the first case above, the dot-notation behavior is similar to that of function `json_value` for a *scalar* value, and it is similar to that of `json_query` for an *object* or *array* value. In the second case, the behavior is similar to that of `json_query` with an array wrapper.

The dot-notation *syntax* is a table alias (mandatory) followed by a dot, that is, a period (`.`), the name of a JSON column, and one or more pairs of the form `. json_field` or `. json_field` followed by `array_step`, where `json_field` is a JSON field name and `array_step` is an array step expression as described in [Basic SQL/JSON Path Expression Syntax](#).

Each `json_field` *must* be a valid SQL identifier,¹ and the column *must* have an `is json` check constraint, which ensures that it contains well-formed JSON data. If either of these rules is not respected then an error is raised at query compile time. (The check constraint must be present to avoid raising an error; however, it need not be active. If you deactivate the constraint then this error is not raised.)

For the dot notation for JSON queries, *unlike the case generally for SQL*, unquoted identifiers (after the column name) are treated *case sensitively*, that is, just as if they were quoted. This is a convenience: you can use JSON field names as identifiers without quoting them. For example, you can write `jcolumn.friends` instead of `jcolumn."friends"`. This also means that if a JSON object is named using uppercase, such as `FRIENDS`, then you must write `jcolumn.FRIENDS`, not `jcolumn.friends`.

Here are some examples of dot notation syntax. All of them refer to JSON column `po_document` of a table that has alias `po`.

- `po.po_document.PONumber` – The value of field `PONumber`.

¹ In particular, this means that you *cannot* use an empty field name (`" "`) with dot-notation syntax.

- `po.po_document.LineItems[1]` – The second element of array `LineItems` (array positions are zero-based).
- `po.po_document.LineItems[*]` – All of the elements of array `LineItems` (* is a wildcard).
- `po.po_document.ShippingInstructions.name` – The value of field `name`, a child of object `ShippingInstructions`.

 **Note:**

- Each component of the dot-notation syntax is limited to a maximum of 128 bytes.
See *Oracle Database SQL Language Reference* for more information about SQL dot-notation syntax and SQL identifiers.
- A simple dot-notation JSON query cannot return a value longer than 4K bytes. If the value surpasses this limit then SQL `NULL` is returned instead. To obtain the actual value, use SQL/JSON function `json_query` or `json_value` instead of dot notation, specifying an appropriate return type with a `RETURNING` clause.
See *Oracle Database SQL Language Reference* for more information about JSON dot-notation syntax.

Matching of a JSON dot-notation expression against JSON data is the same as matching of a SQL/JSON path expression, including the relaxation to allow implied array iteration (see [SQL/JSON Path Expression Syntax Relaxation](#)). The JSON column of a dot-notation expression corresponds to the context item of a path expression, and each identifier used in the dot notation corresponds to an identifier used in a path expression.

For example, if JSON column `jcolumn` corresponds to the path-expression context item, then the expression `jcolumn.friends` corresponds to path expression `$.friends`, and `jcolumn.friends.name` corresponds to path expression `$.friends.name`.

For the latter example, the context item could be an object or an array of objects. If it is an array of objects then each of the objects in the array is matched for a field `friends`. The value of field `friends` can itself be an object or an array of objects. In the latter case, the first object in the array is used.

 **Note:**

Other than (1) the *implied* use of a wildcard for array elements (see [SQL/JSON Path Expression Syntax Relaxation](#)) and (2) the explicit use of a wildcard between array brackets ([*]), you *cannot* use wildcards in a path expression when you use the dot-notation syntax. This is because an asterisk (*) is not a valid *SQL identifier*.

For example, this raises a syntax error:
`mytable.mycolumn.object1.*.object2.`

Dot-notation syntax is a handy alternative to using simple path expressions; it is not a replacement for using path expressions in general.

[Example 12-1](#) shows equivalent dot-notation and `json_value` queries. Given the data from [Example 4-2](#), each of the queries returns the string "1600", a `VARCHAR2` value representing the JSON number 1600.

[Example 12-2](#) shows equivalent dot-notation and `json_query` queries. Each query in the first pair returns (a `VARCHAR2` value representing) a JSON array of phone objects. Each query in the second pair returns (a `VARCHAR2` value representing) an array of phone types, just as in [Example 17-1](#).

 **See Also:**

Oracle Database SQL Language Reference for information about dot notation used for SQL object and object attribute access (object access expressions)

Example 12-1 JSON Dot-Notation Query Compared With JSON_VALUE

```
SELECT po.po_document.PONumber FROM j_purchaseorder po;

SELECT json_value(po_document, '$.PONumber') FROM j_purchaseorder;
```

Example 12-2 JSON Dot-Notation Query Compared With JSON_QUERY

```
SELECT po.po_document.ShippingInstructions.Phone FROM j_purchaseorder
po;

SELECT json_query(po_document, '$.ShippingInstructions.Phone')
FROM j_purchaseorder;

SELECT po.po_document.ShippingInstructions.Phone.type FROM
j_purchaseorder po;

SELECT json_query(po_document, '$.ShippingInstructions.Phone.type' WITH
```



```
WRAPPER)  
FROM j_purchaseorder;
```

Related Topics

- [Overview of SQL/JSON Path Expressions](#)
Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.
- [Creating a Table With a JSON Column](#)
You can create a table that has JSON columns. You use SQL condition `is json` as a check constraint to ensure that data inserted into a column is (well-formed) JSON data. Oracle recommends that you *always* use an `is_json` check constraint when you create a column intended for JSON data.
- [SQL/JSON Function JSON_TABLE](#)
SQL/JSON function `json_table` projects specific JSON data to columns of various SQL data types. You use it to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.

13

SQL/JSON Path Expressions

Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.

- [Overview of SQL/JSON Path Expressions](#)
Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.
- [SQL/JSON Path Expression Syntax](#)
SQL/JSON path expressions are matched by SQL/JSON functions and conditions against JSON data, to select portions of it. Path expressions can use wildcards and array ranges. Matching is case-sensitive.
- [SQL/JSON Path Expression Item Methods](#)
The Oracle item methods available for a SQL/JSON path expression are described.
- [ISO 8601 Date and Time Support](#)
International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. Oracle Database supports many of the ISO 8601 date and time formats.
- [Types in Comparisons](#)
Comparisons in SQL/JSON path-expression filter conditions are statically typed at compile time. If the effective types of the operands of a comparison are not known to be the same then an attempt is sometimes made to reconcile them by type-casting.

13.1 Overview of SQL/JSON Path Expressions

Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.

JSON is a notation for JavaScript values. When JSON data is stored in the database you can query it using path expressions that are somewhat analogous to XQuery or XPath expressions for XML data. Similar to the way that SQL/XML allows SQL access to XML data using XQuery expressions, Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.

SQL/JSON path expressions have a simple syntax. A path expression selects zero or more JSON values that match, or satisfy, it.

SQL/JSON condition `json_exists` returns true if at least one value matches, and false if no value matches. If a single value matches, then SQL/JSON function `json_value` returns that value if it is scalar and raises an error if it is non-scalar. If no value matches the path expression then `json_value` returns SQL NULL.

SQL/JSON function `json_query` returns all of the matching values, that is, it can return multiple values. You can think of this behavior as returning a sequence of values, as in XQuery, or you can think of it as returning multiple values. (No user-visible sequence is manifested.)

In all cases, path-expression matching attempts to match each *step* of the path expression, in turn. If matching any step fails then no attempt is made to match the subsequent steps, and matching of the path expression fails. If matching each step succeeds then matching of the path expression succeeds.

Related Topics

- [SQL/JSON Path Expression Syntax](#)
SQL/JSON path expressions are matched by SQL/JSON functions and conditions against JSON data, to select portions of it. Path expressions can use wildcards and array ranges. Matching is case-sensitive.

13.2 SQL/JSON Path Expression Syntax

SQL/JSON path expressions are matched by SQL/JSON functions and conditions against JSON data, to select portions of it. Path expressions can use wildcards and array ranges. Matching is case-sensitive.

You pass a SQL/JSON path expression and some JSON data to a SQL/JSON function or condition. The path expression is matched against the data, and the matching data is processed by the particular SQL/JSON function or condition. You can think of this matching process in terms of the path expression *returning* the matched data to the function or condition.

- [Basic SQL/JSON Path Expression Syntax](#)
The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (\$) followed by zero or more object, array, and descendant steps, each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.
- [SQL/JSON Path Expression Syntax Relaxation](#)
The basic SQL/JSON path-expression syntax is relaxed to allow implicit array wrapping and unwrapping. This means that you need not change a path expression in your code if your data evolves to replace a JSON value with an array of such values, or vice versa. Examples are provided.

Related Topics

- [About Strict and Lax JSON Syntax](#)
The Oracle default syntax for JSON is lax. In particular: it reflects the JavaScript syntax for object fields; the Boolean and null values are not case-sensitive; and it is more permissive with respect to numerals, whitespace, and escaping of Unicode characters.
- [Diagrams for Basic SQL/JSON Path Expression Syntax](#)
Syntax diagrams and corresponding Backus-Naur Form (BNF) syntax descriptions are presented for the basic SQL/JSON path expression syntax.

13.2.1 Basic SQL/JSON Path Expression Syntax

The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (\$) followed by zero or more object, array, and descendant steps,

each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.

However, this basic syntax is extended by relaxing the matching of arrays and non-arrays against non-array and array patterns, respectively — see [SQL/JSON Path Expression Syntax Relaxation](#).

Matching of data against SQL/JSON path expressions is case-sensitive.

- A SQL/JSON **basic path expression** (also called just a *path expression* here) is an *absolute path expression* or a *relative path expression*.
- An **absolute path expression** begins with a dollar sign (\$), which represents the path-expression **context item**, that is, the JSON data to be matched. That data is the result of evaluating a SQL expression that is passed as argument to the SQL/JSON function. The dollar sign is followed by zero or more *nonfunction steps*, followed by an optional *function step*.
- A **relative path expression** is an at sign (@) followed by zero or more *nonfunction steps*, followed by an optional *function step*. It has the same syntax as an *absolute path expression*, except that it uses an at sign instead of a dollar sign (\$).

A relative path expression is used inside a *filter expression* (*filter*, for short). The at sign represents the path-expression **current filter item**, that is, the JSON data that matches the part of the (surrounding) path expression that precedes the filter containing the relative path expression. A relative path expression is matched against the current filter item in the same way that an absolute path expression is matched against the context item.

- A **nonfunction step** is an *object step*, an *array step*, or a *descendant step*, followed by an optional *filter expression*.
- A single **function step** is *optional* in a *basic path expression* (absolute or a relative). If present, it is the last step of the path expression. It is a period (.), sometimes read as "dot", followed by a SQL/JSON **item method**, followed by a left parenthesis (()) and then a right parenthesis ()). The parentheses can have whitespace between them (such whitespace is insignificant).

The item method is applied to the data that is targeted by the rest of the same path expression, which *precedes* the function step. The item method is used to transform that data. The SQL function or condition that is passed the path expression as argument uses the transformed data in place of the targeted data.

- An **object step** is a period (.), followed by an object field name or an asterisk (*) wildcard, which stands for (the values of) *all* fields. A field name can be *empty*, in which case it *must* be written as "" (no intervening whitespace). A nonempty field name must start with an uppercase or lowercase letter A to Z and contain only such letters or decimal digits (0-9), or else it must be enclosed in double quotation marks (").

An object step returns the *value* of the field that is specified. If a wildcard is used for the field then the step returns the values of *all* fields, in no special order.

- An **array step** is a left bracket ([]) followed by *either* an asterisk (*) wildcard, which stands for *all* array elements, *or* one or more specific array indexes or range specifications separated by commas (,), followed by a right bracket (]). An error is raised if you use both an asterisk and either an array index or a range specification.

In a path expression, **array indexing** is zero-based (0, 1, 2,...), as in the JavaScript convention for arrays. A **range specification** has the form *N* to *M*,

where N and M are array indexes and N is strictly less than M .¹ An error is raised at query compilation time if N is not less than M .

When indexes or range specifications are used, the array elements they collectively specify must be specified in ascending order, without repetitions, or else a compile-time error is raised. For example, an error is raised for each of [3, 1 to 4], [4, 2], [2, 3 to 3], and [2, 3, 3]. Errors are raised on the first two because the order is not ascending, Errors are raised on the last two because of the repetition of array index 3 (which indexes the fourth array element, because of zero-based indexing).

Similarly, the elements in the array value that results from matching are in ascending order, with no repetitions. If an asterisk is used in the path expression then all of the array elements are returned, in array order.

- A **descendant** step is two consecutive periods (`..`), sometimes read as "dot dot", followed by a field name (which has the same syntax as for an *object step*).

It *descends recursively* into the objects or arrays that match the step immediately preceding it (or into the context item if there is no preceding step).

At each descendant level, for each object and for each array element that is an object, it gathers the values of all fields that have the specified name. It returns all of the gathered field values.

For example, consider this query and data:

```
json_query(some_json_column, '$..z' WITH ARRAY WRAPPER)
```

```
{ "a" : { "b" : { "z" : 1 },
          "c" : [ 5, { "z" : 2 } ],
          "z" : 3 }
  "z" : 4 }
```

The query returns an array, such as [1, 2, 3], whose elements are 1, 2, and 3. It gathers the value of each field `z` within the step that immediately precedes the dot dot (`..`), which is field `a`. The topmost field `z`, with value 4, is *not* matched because it is not within the value of field `a`.

The value of field `a` is an object, which is descended into.

- It has a field `z`, whose value (3) is gathered. It also has a field `b` whose value is an object, which is descended into to gather the value of its field `z`, which is 1.
- It also has a field `c` whose value is an array, which has an element that is an object with a field `z`, whose value (2) is gathered.

The JSON values gathered are thus 3, 1, and 2. They are wrapped in an array, in an undefined order. One of the possible return values is [1, 2, 3].

- A **filter expression** (**filter**, for short) is a question mark (?) followed by a *filter condition* enclosed in parentheses (`()`). A filter is satisfied if its condition is satisfied, that is, returns true.
- A **filter condition** applies a predicate (Boolean function) to its arguments and is one of the following, where each of *cond*, *cond1*, and *cond2* stands for a filter condition.

¹ The `to` in a range specification is sometimes informally called the array *slice* operator.

- (*cond*): Parentheses are used for *grouping*, separating filter condition *cond* as a unit from other filter conditions that may precede or follow it.
- *cond1* && *cond2*: The *conjunction (and)* of *cond1* and *cond2*, requiring that both be satisfied.
- *cond1* || *cond2*: The inclusive *disjunction (or)* of *cond1* and *cond2*, requiring that *cond1*, *cond2*, or both, be satisfied.
- ! (*cond*): The *negation* of *cond*, meaning that *cond* must *not* be satisfied.
- **exists** (, followed by a *relative path expression*, followed by): The condition that the targeted data *exists* (is present).
- A **comparison**, which is one of the following:
 - * A *relative path expression*, followed by a *comparison predicate*, followed by either a JSON scalar value or a *SQL/JSON variable*.
 - * Either a JSON scalar value or a *SQL/JSON variable*, followed by a *comparison predicate*, followed by a *relative path expression*.
 - * A JSON scalar value, followed by a *comparison predicate*, followed by another JSON scalar value.
 - * A *relative path expression*, followed by *has substring*, *starts with*, *like*, *like_regex*, or *eq_regex*, followed by either a JSON string or a *SQL/JSON variable* that is bound to a SQL string (which is automatically converted from the database character set to UTF8).
 - * **has substring** means that the matching data value has the specified string as a *substring*.
 - * **starts with** means that the matching data value has the specified string as a *prefix*.
 - * **like** means that the JSON string data value matches the specified string, which is interpreted as a SQL `LIKE` pattern that uses SQL `LIKE4` character-set semantics. A percent sign (%) in the pattern matches zero or more characters. An underscore (_) matches a single character.

 **Note:**

Unlike the case for SQL `LIKE`, there is *no escape character* for path-expression predicate `like`. Also, Oracle recommends that you avoid using character ```, GRAVE ACCENT (U+0060), in your `like` patterns — that character, also known sometimes as backquote or backtick, is *reserved for future use*.

- * **like_regex** means that the JSON string data value matches the specified string, which is interpreted as a SQL `REGEXP LIKE` *regular expression* pattern that uses SQL `LIKE4` character-set semantics.
`like_regex` is *exceptional* among the pattern-matching comparisons, in that its pattern matches the empty JSON string ("").
- * **eq_regex** is just like `like_regex`, except for these two differences:

- * `eq_regex` matches its regular expression pattern against the entire JSON string data value — *the full string must match* the pattern for the comparison to be satisfied. `like_regex` is satisfied if any portion of the JSON string matches the pattern.
- * The `eq_regex` pattern does not match the empty JSON string ("").

For all of these predicates, a pattern that is the empty string ("") matches data that is the empty string. And for all except `like_regex`, a pattern that is a nonempty string does *not* match data that is the empty string. For `like_regex` a nonempty pattern does match empty-string data.

- * A *relative path expression*, followed by `in`, followed by a *value list*, meaning that the value is one of those in the *value list*.

A **comparison predicate** is `==`, `<>`, `!=`², `<`, `<=`, `>=`, or `>`, meaning equals, does not equal, is less than, is less than or equal to, is greater than or equal to, and is greater than, respectively.

A **SQL/JSON variable** is a dollar sign (\$) followed by the name of a SQL identifier that is bound in a `PASSING` clause for `json_exists`.

- A **value list** is (, followed by a list of one or more scalar values and *SQL/JSON variables* separated by commas (,), followed by).

The predicates that you can use in filter conditions are thus `&&`, `||`, `!`, `exists`, `==`, `<>`, `!=`, `<`, `<=`, `>=`, `>`, and `in`.

As an example, the filter condition `(a || b) && (!(c) || d < 42)` is satisfied if both of the following criteria are met:

- At least one of the filter conditions `a` and `b` is satisfied: `(a || b)`.
- Filter condition `c` is *not* satisfied or the number `d` is less than or equal to 42, or both are true: `!(c) || d < 42)`.

Comparison predicate `!` has precedence over `&&`, which has precedence over `||`. You can always use parentheses to control grouping.

Without parentheses for grouping, the preceding example would be `a || b && !(c) || d < 42`, which would be satisfied if at least one of the following criteria is met:

- Condition `b && !(c)` is satisfied, which means that each of the conditions `b` and `!(c)` is satisfied (which in turn means that condition `c` is not satisfied).
- Condition `a` is satisfied.
- Condition `d < 42` is satisfied.

At least one side of a comparison must *not* be a SQL/JSON variable. The default *type* for a comparison is defined at compile time, based on the type(s) for the non-variable side(s). You can use a type-specifying *item method* to override this default with a different type. The type of your matching data is automatically converted, for the comparison, to fit the determined type (default or specified by item method). For example, `$.a > 5` imposes numerical comparison because 5 is a number, `$.a > "5"` imposes string comparison because "5" is a string.

² `!=` is an Oracle alias for the SQL/JSON standard comparison predicate `<>`.

 **Note:**

For powerful full-text search, use Oracle SQL function `json_textcontains`, which requires that you create a JSON search index. As a less-powerful alternative, if you do not create a JSON search index, and you just want simple string pattern-matching in a filter condition, you can use any of the pattern-matching comparisons: `has substring`, `starts with`, `like`, `like_regex`, or `eq_regex`.

Here are some examples of path expressions, with their meanings spelled out in detail.

- `$` – The context item.
- `$.friends` – The value of field `friends` of a context-item object. The dot (`.`) immediately after the dollar sign (`$`) indicates that the context item is a JSON *object*.
- `$.friends[0]` – An object that is the first element of an array that is the value of field `friends` of a context-item object. The bracket notation indicates that the value of field `friends` is an *array*.
- `$.friends[0].name` – Value of field `name` of an object that is the first element of an array that is the value of field `friends` of a context-item object. The second dot (`.`) indicates that the first element of array `friends` is an object (with a `name` field).
- `$.friends[*].name` – Value of field `name` of *each* object in an array that is the value of field `friends` of a context-item object.
- `$.*[*].name` – Field `name` values for each object in an array value of a field of a context-item object.
- `$.friends[3, 8 to 10, 12]` – The fourth, ninth through eleventh, and thirteenth elements of an array `friends` (field of a context-item object). The elements must be specified in *ascending order*, and they are returned in that order: fourth, ninth, tenth, eleventh, thirteenth.
- `$.friends[3].cars` – The value of field `cars` of an object that is the fourth element of an array `friends`. The dot (`.`) indicates that the fourth element is an object (with a `cars` field).
- `$.friends[3].*` – The values of *all* of the fields of an object that is the fourth element of an array `friends`.
- `$.friends[3].cars[0].year` – The value of field `year` of an object that is the first element of an array that is the value of field `cars` of an object that is the fourth element of an array `friends`.
- `$.friends[3].cars[0]?(@.year > 2016)` – The first object of an array `cars` (field of an object that is the fourth element of an array `friends`), *provided that* the value of its field `year` is, or can be converted to, a number greater than 2016. A `year` value such as "2017" is converted to the number 2017, which satisfies the test. A `year` value such as "recent" fails the test — no match.
- `$.friends[3].cars[0]?(@.year.number() > 2016)` – Same as the previous. Item method `number()` allows only a number or a string value that can be converted to

a number, and that behavior is already provided by numeric comparison predicate `>`.

- `$.friends[3].cars[0]?(@.year.numberOnly() > 2016)` – Same as the previous, but only if the `year` value is a number. Item method `numberOnly()` excludes a car with a `year` value that is a string numeral, such as "2017".
- `$.friends[3]?(@.addresses.city == "San Francisco")` – An object that is the fourth element of an array `friends`, provided that it has an `addresses` field whose value is an object with a field `city` whose value is the string "San Francisco".
- `$.friends[*].addresses?(@city starts with "San ").zip` – Zip codes of all addresses of `friends`, where the name of the address `city` starts with "San ". (In this case the filter is not the last path step.)
- `$.zip` – All values of a `zip` field, anywhere, at any level.
- `$.friends[3]?(@.addresses.city == "San Francisco" && @.addresses.state == "Nevada")` – Objects that are the fourth element of an array `friends`, provided that there is a match for an address with a `city` of "San Francisco" and there is a match for an address with a `state` of "Nevada".

Note: The filter conditions in the conjunction do *not* necessarily apply to the same object — the filter tests for the existence of an object with `city` San Francisco and for the existence of an object with `state` Nevada. It does *not* test for the existence of an object with both `city` San Francisco and `state` Nevada. See [Using Filters with JSON_EXISTS](#).

- `$.friends[3].addresses?(@.city == "San Francisco" && @.state == "Nevada")` – An object that is the fourth element of array `friends`, provided that object has a match for `city` of "San Francisco" and a match for `state` of "Nevada".

Unlike the preceding example, in this case the filter conditions in the conjunction, for fields `city` and `state`, apply to the *same* `addresses` object. The filter applies to a given `addresses` object, which is outside it.

Related Topics

- [Using Filters with JSON_EXISTS](#)
You can use SQL/JSON condition `json_exists` with a path expression that has one or more filter expressions, to select documents that contain matching data. Filters let you test for the existence of documents that have particular fields that satisfy various conditions.
- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no **RETURNING** clause) are described here.
- [SQL/JSON Path Expression Item Methods](#)
The Oracle item methods available for a SQL/JSON path expression are described.
- [SQL/JSON Path Expression Syntax Relaxation](#)
The basic SQL/JSON path-expression syntax is relaxed to allow implicit array wrapping and unwrapping. This means that you need not change a path expression in your code if your data evolves to replace a JSON value with an array of such values, or vice versa. Examples are provided.

- [Diagrams for Basic SQL/JSON Path Expression Syntax](#)
Syntax diagrams and corresponding Backus-Naur Form (BNF) syntax descriptions are presented for the basic SQL/JSON path expression syntax.
- [Wrapper Clause for SQL/JSON Query Functions JSON_QUERY and JSON_TABLE](#)
SQL/JSON query functions `json_query` and `json_table` accept an optional wrapper clause, which specifies the form of the value returned by `json_query` or used for the data in a `json_table` column. This clause and the default behavior (no wrapper clause) are described here. Examples are provided.
- [ISO 8601 Date and Time Support](#)
International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. Oracle Database supports many of the ISO 8601 date and time formats.

 **See Also:**

- ISO 8601 for information about the ISO date formats
- *Oracle Database SQL Language Reference* for information about SQL condition `REGEXP LIKE`
- *Oracle Database SQL Language Reference* for information about SQL condition `LIKE` and `LIKE4` character-set semantics

13.2.2 SQL/JSON Path Expression Syntax Relaxation

The basic SQL/JSON path-expression syntax is relaxed to allow implicit array wrapping and unwrapping. This means that you need not change a path expression in your code if your data evolves to replace a JSON value with an array of such values, or vice versa. Examples are provided.

[Basic SQL/JSON Path Expression Syntax](#) defines the basic SQL/JSON path-expression syntax. The actual path expression syntax supported relaxes that definition as follows:

- If a path-expression step targets (expects) an array but the actual data presents no array then the data is implicitly wrapped in an array.
- If a path-expression step targets (expects) a non-array but the actual data presents an array then the array is implicitly unwrapped.

This relaxation allows for the following abbreviation: `[*]` can be elided whenever it precedes the object accessor, `.`, followed by an object field name, with no change in effect. The reverse is also true: `[*]` can always be inserted in front of the object accessor, `.`, with no change in effect.

This means that the object step `[*].prop`, which stands for the value of field `prop` of each element of a given array of objects, can be abbreviated as `.prop`, and the object step `.prop`, which looks as though it stands for the `prop` value of a single object, stands also for the `prop` value of each element of an array to which the object accessor is applied.

This is an important feature, because it means that you need not change a path expression in your code if your data evolves to replace a given JSON value with an array of such values, or vice versa.

For example, if your data originally contains objects that have field `Phone` whose value is a single object with fields `type` and `number`, the path expression `$.Phone.number`, which matches a single phone number, can still be used if the data evolves to represent an array of phones. Path expression `$.Phone.number` matches either a single phone object, selecting its number, or an array of phone objects, selecting the number of each.

Similarly, if your data mixes both kinds of representation — there are some data entries that use a single phone object and some that use an array of phone objects, or even some entries that use both — you can use the same path expression to access the phone information from these different kinds of entry.

Here are some example path expressions from section [Basic SQL/JSON Path Expression Syntax](#), together with an explanation of equivalences.

- `$.friends` – The value of field `friends` of *either*:
 - The (single) context-item object.
 - (equivalent to `$(*).friends`) Each object in the context-item array.
- `$.friends[0].name` – Value of field `name` for *any* of these objects:
 - The first element of the array that is the value of field `friends` of the context-item object.
 - (equivalent to `$.friends.name`) The value of field `friends` of the context-item object.
 - (equivalent to `$(*).friends.name`) The value of field `friends` of each object in the context-item array.
 - (equivalent to `$(*).friends[0].name`) The first element of each array that is the value of field `friends` of each object in the context-item array.

The context item can be an object or an array of objects. In the latter case, each object in the array is matched for a field `friends`.

The value of field `friends` can be an object or an array of objects. In the latter case, the first object in the array is used.

- `$.*[*].name` – Value of field `name` for *any* of these objects:
 - An element of an array value of a field of the context-item object.
 - (equivalent to `$.*.name`) The value of a field of the context-item object.
 - (equivalent to `$(*)*.name`) The value of a field of an object in the context-item array.
 - (equivalent to `$(*).*[*].name`) Each object in an array value of a field of an object in the context-item array.

Related Topics

- [Basic SQL/JSON Path Expression Syntax](#)
The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (`$`) followed by zero or more object, array, and descendant

steps, each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.

13.3 SQL/JSON Path Expression Item Methods

The Oracle item methods available for a SQL/JSON path expression are described.

An item method is applied to the JSON data that is targeted by (the rest of) the path expression terminated by that method. The method is used to transform that data. The SQL function or condition that is passed the path expression uses the transformed data in place of the targeted data. In some cases the application of an item method acts as a filter, removing the targeted data from the result set.

If an item-method conversion fails for any reason, such as its argument being of the wrong type, then the path cannot be matched (it refers to no data), and *no error is raised*. In particular, this means that such an error is *not* handled by an error clause in the SQL/JSON function or condition to which the path expression is passed.

Application of an Item Method to an Array

With the exception of item methods `size()` and `type()`, if an array is targeted by an item method then the method is *applied to each of the array elements*, not to the array itself. For example, `$.a.method()` applies item-method `method()` to each element of array `a`, to convert it and use it in place of the array. (This is similar, in effect, to the implied unwrapping of an array when a non-array is expected for an object step.) The resulting set of matches includes the converted array elements, not the targeted array.

For `json_query` or a `json_table` column expression with `json_query` semantics you can use a wrapper clause to capture all of the converted array-element values as an array. For example, this query:

```
SELECT json_query(['"alpha"', 42, '"10.4"'], '$[*].stringOnly()') WITH
ARRAY WRAPPER) FROM dual;
```

returns this array (as a `VARCHAR2` value): `["alpha", "10.4"]`.

Item methods `size()` and `type()` are *exceptional* in this regard. When applied to an array they treat it as such, instead of acting on its elements. For example:

```
SELECT json_value('[19, "Oracle", {"a":1},[1,2,3]]', '$.type()') FROM
dual;
```

returns the single `VARCHAR2` value `'array'`.

Item-Method Descriptions

In the item-method descriptions here, in some cases a targeted JSON value is said to be *interpreted as a value of a given SQL data type*. This means that it is handled as if it were controlled by a `json_value RETURNING` clause with that SQL data type.

For example, item-method `string()` interprets its target as would `json_value` with clause `RETURNING VARCHAR2(4000)`. A Boolean value is thus treated by `string()` as `"true"` or `"false"`; a null value is treated as `"null"`; and a number is represented in a canonical string form.

Because some item methods interpret the targeted JSON data as if it were of a SQL data type, they can be used with `json_value` in place of a `RETURNING` clause, and they can be used with `json_table` in place of a column type specification. That is, the *item methods can be used to specify the SQL data type* for the extracted JSON data.

You can also use such item methods *together with* a `json_value` `RETURNING` clause or a `json_table` column type specification. What happens if the SQL data type to use for extracted JSON data is controlled by *both* an item method and either a `json_value` `RETURNING` clause or a `json_table` column type?

- If the two data types are compatible then the data type for the `RETURNING` clause or the column is used. For these purposes, `VARCHAR2` is compatible with both `VARCHAR2` and `CLOB`.
- If the data types are incompatible then a static, compile-time *error* is raised.

Table 13-1 Compatibility of Type-Conversion Item Methods and RETURNING Types

Item Method	Compatible RETURNING Clause Data Types
<code>string()</code> or <code>stringOnly()</code>	<code>VARCHAR2</code> or <code>CLOB</code> , except that <code>string()</code> returns SQL <code>NULL</code> for a JSON <code>null</code> value
<code>number()</code> or <code>numberOnly()</code>	<code>NUMBER</code>
<code>date()</code>	<code>DATE</code>
<code>timestamp()</code>	<code>TIMESTAMP</code>
<code>boolean()</code> or <code>booleanOnly()</code>	<code>VARCHAR2</code>

Using a `RETURNING` clause or a column specification, you can specify a length for character data and a precision and scale for numerical data. This lets you assign a more precise SQL data type for extraction than what is provided by an item method for target-data comparison purposes.

For example, if you use item method `string()` and `RETURNING VARCHAR2(150)` then the data type of the returned data is `VARCHAR2(150)`, not `VARCHAR2(4000)`.

The data-type conversion methods with “only” in their name are the same as the corresponding methods with names without “only”, except that the former convert *only* JSON values that are of the given type (e.g., `number`) to the related SQL data type (e.g. `NUMBER`). The methods without “only” in the name allow conversion, when possible, of *any* JSON value to the given SQL data type. (When an “only” method targets an array, the conversion applies to each array element, as usual.)

- `abs()`: The absolute value of the targeted JSON number. Corresponds to the use of SQL function `ABS`.
- `boolean()`: A SQL `VARCHAR2(20)` interpretation of the targeted JSON value.
- `booleanOnly()`: A SQL `VARCHAR2(20)` interpretation of the targeted JSON data, but only if it is a JSON Boolean value; otherwise, there is no match. Acts as a filter, allowing matches only for JSON Boolean values.
- `ceiling()`: The targeted JSON number, rounded up to the nearest integer. Corresponds to the use of SQL function `CEIL`.
- `date()`: A SQL date interpretation of the targeted JSON string. The targeted string data must be in one of the supported ISO 8601 formats; otherwise, there is no match.

- **double()**: A SQL `BINARY_DOUBLE` interpretation of the targeted JSON string or number.
- **floor()**: The targeted JSON number, rounded down to the nearest integer. Corresponds to the use of SQL function `FLOOR`.
- **length()**: The number of characters in the targeted JSON string, interpreted as a SQL `NUMBER`.
- **lower()**: The lowercase string that corresponds to the characters in the targeted JSON string.
- **number()**: A SQL `NUMBER` interpretation of the targeted JSON string or number.
- **numberOnly()**: A SQL `NUMBER` interpretation of the targeted JSON data, but only if it is a JSON number; otherwise, there is no match. Acts as a filter, allowing matches only for JSON numbers.
- **size()**: The number of elements in an array, or 1 for a scalar or an object. This item method can be used with `json_query`, in addition to `json_value` and `json_table`. If applied to data that is an array, no implicit iteration over the array elements occurs: the resulting value is just the number of array elements. (This is an exception to the rule of implicit iteration.)
- **string()**: A SQL `VARCHAR2(4000)` interpretation of the targeted scalar JSON value.
- **stringOnly()**: A SQL `VARCHAR2(4000)` interpretation of the targeted scalar JSON value, but only if it is a JSON string; otherwise, there is no match. Acts as a filter, allowing matches only for JSON strings.
- **timestamp()**: A SQL `TIMESTAMP` interpretation of the targeted JSON string. The targeted string data must be in one of the supported ISO 8601 formats; otherwise, there is no match.³
- **type()**: The name of the JSON data type of the targeted data, interpreted as a SQL `VARCHAR2(20)` value. This item method can be used with `json_query`, in addition to `json_value` and `json_table`. If applied to data that is an array, no implicit iteration over the array elements occurs: the resulting value is "array". (This is an exception to the rule of implicit iteration.)
 - "null" for a value of null.
 - "boolean" for a value of true or false.
 - "number" for a number.
 - "string" for a string.
 - "array" for an array.
 - "object" for an object.
- **upper()**: The uppercase string that corresponds to the characters in the targeted JSON string.

Item methods `boolean()`, `booleanOnly()`, `date()`, `length()`, `lower()`, `number()`, `numberOnly()`, `string()`, `stringOnly()`, `timestamp()`, and `upper()` are Oracle extensions to the SQL/JSON standard. The other item methods, `abs()`, `ceiling()`, `double()`, `floor()`, `size()`, and `type()` are part of the standard.

³ Applying item method `timestamp()` to a supported ISO 8601 string `<ISO-STRING>` has the effect of SQL `sys_extract_utc(to_utc_timestamp_tz(<ISO-STRING>))`.

Related Topics

- [Basic SQL/JSON Path Expression Syntax](#)
The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (\$) followed by zero or more object, array, and descendant steps, each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.
- [ISO 8601 Date and Time Support](#)
International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. Oracle Database supports many of the ISO 8601 date and time formats.
- [Types in Comparisons](#)
Comparisons in SQL/JSON path-expression filter conditions are statically typed at compile time. If the effective types of the operands of a comparison are not known to be the same then an attempt is sometimes made to reconcile them by type-casting.
- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no **RETURNING** clause) are described here.
- [SQL/JSON Function JSON_VALUE](#)
SQL/JSON function `json_value` selects JSON data and returns a SQL scalar or an instance of a user-defined SQL object type or SQL collection type (varray, nested table).
- [SQL/JSON Function JSON_TABLE](#)
SQL/JSON function `json_table` projects specific JSON data to columns of various SQL data types. You use it to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.
- [Wrapper Clause for SQL/JSON Query Functions JSON_QUERY and JSON_TABLE](#)
SQL/JSON query functions `json_query` and `json_table` accept an optional wrapper clause, which specifies the form of the value returned by `json_query` or used for the data in a `json_table` column. This clause and the default behavior (no wrapper clause) are described here. Examples are provided.

13.4 ISO 8601 Date and Time Support

International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. Oracle Database supports many of the ISO 8601 date and time formats.

International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. You can manipulate strings that are in the most common ISO 8601 date and time formats as proper Oracle Database date and time values. The ISO 8601 formats that are supported are essentially those that are numeric-only, language-neutral, and unambiguous.

This is the allowed syntax for dates and times:

- Date (only): `YYYY-MM-DD`
- Date with time: `YYYY-MM-DDThh:mm:ss[.s[s[s[s[s[s]]]]][Z|(+|-)hh:mm]`

where:

- **YYYY** specifies the *year*, as four decimal digits.
- **MM** specifies the *month*, as two decimal digits, 00 to 12.
- **DD** specifies the *day*, as two decimal digits, 00 to 31.
- **hh** specifies the *hour*, as two decimal digits, 00 to 23.
- **mm** specifies the *minutes*, as two decimal digits, 00 to 59.
- **ss[.s[s[s[s[s]]]]]** specifies the *seconds*, as two decimal digits, 00 to 59, optionally followed by a decimal point and 1 to 6 decimal digits (representing the fractional part of a second).
- **z** specifies *UTC* time (time zone 0). (It can also be specified by +00:00, but not by -00:00.)
- **(+|-)hh:mm** specifies the time-zone as *difference from UTC*. (One of + or - is required.)

For a time value, the time-zone part is optional. If it is absent then UTC time is assumed.

No other ISO 8601 date-time syntax is supported. In particular:

- Negative dates (dates prior to year 1 BCE), which begin with a hyphen (e.g. -2018-10-26T21:32:52), are not supported.
- Hyphen and colon separators are required: so-called “basic” format, YYYYMMDDThhmmss, is not supported.
- Ordinal dates (year plus day of year, calendar week plus day number) are not supported.
- Using more than four digits for the year is not supported.

Supported dates and times include the following:

- 2018-10-26T21:32:52
- 2018-10-26T21:32:52+02:00
- 2018-10-26T19:32:52Z
- 2018-10-26T19:32:52+00:00
- 2018-10-26T21:32:52.12679

Unsupported dates and times include the following:

- 2018-10-26T21:32 (if a time is specified then all of its parts must be present)
- 2018-10-26T**25**:32:52+02:00 (the hours part, 25, is out of range)
- 18-10-26T21:32 (the year is not specified fully)

See Also:

- ISO 8601 standard
- [ISO 8601 at Wikipedia](#)

13.5 Types in Comparisons

Comparisons in SQL/JSON path-expression filter conditions are statically typed at compile time. If the effective types of the operands of a comparison are not known to be the same then an attempt is sometimes made to reconcile them by type-casting.

A SQL/JSON path expression targets JSON data, so the operands of a comparison are JSON values. Type comparison of JSON values is straightforward: JSON data types string, number, null, object, and array are mutually exclusive and incomparable.

But comparison operands are sometimes *interpreted* (essentially cast) as values of SQL data types. This is the case, for example, when some item methods, such as `number()`, are used. This section addresses the type-checking of such effective values.

You can prevent such type-casting by explicitly using one of the “only” item methods. For example, applying method `numberOnly()` prevents implicit type-casting to a number.

SQL is a statically typed language; types are determined at compile time. The same applies to SQL/JSON path expressions, and in particular to comparisons in filter conditions. This means that you get the same result for a query regardless of how it is evaluated — whether functionally or using features such as indexes, materialized views, and In-Memory scans.

To realize this:

- If the types of both operands are *known* and they are the *same* then type-checking is satisfied.
- If the types of both operands are *unknown* then a compile-time error is raised.
- If the type of one operand is known and the other is unknown then the latter operand is cast to the type of the former.

For example, in `$.a?(@.b.c == 3)` the type of `$.a.b.c` is unknown at compile time. The path expression is compiled as `$.a?(@.b.c.number() == 3)`. At runtime an attempt is thus made to cast the data that matches `$.a.b.c` to a number. A string value `"3"` would be cast to the number `3`, satisfying the comparison.⁴

- If the types of both operands are *known* and they are *not* the same then an attempt is made to cast the type of one to the type of the other. Details are presented below.

Comparison operands used in the following combinations are *not* reconciled; a *compile-time error* is raised.

- Number compared with any non-number type other than double.
- Double compared with any non-double type other than number.
- String compared with double.
- Boolean compared with any non-Boolean type, except for comparison with a string value of `'true'` or `'false'` (or any letter-case variants, such as `'tRUe'`).

⁴ To prevent such casting here, you can explicitly apply item method `numberOnly()`: `$.a?(@.b.c.numberOnly() == 3)`. Data with a string value `"3"` would simply not match; it would be filtered out.

- Date or timestamp compared with string, unless the string has a supported ISO 8601 format.
- Date compared with any non-date type other than string.
- Timestamp (with or without time zone) compared with any non-timestamp type other than string.
- JSON `null` compared with any type other than JSON `null`.

An attempt is made to reconcile comparison operands used in the following combinations, by type-casting. You can think of a type-casting item method being implicitly applied to one of the operands in order to make it type-compatible with the other operand.

- Number compared with double — `double()` is implicitly applied to the number to make it a double value.
- String value of `'true'` or `'false'` (or any letter-case variants, such as `'tRUe'`) compared with Boolean — `boolean()` is implicitly applied to the string to make it a Boolean value.
- String in a supported ISO 8601 format compared with date — `date()` is implicitly applied to the string to make it a date value. (Any time-zone component present is removed.)
- String in a supported ISO 8601 format compared with timestamp (with or without) time zone — `timestamp()` is implicitly applied to the string to make it a timestamp value. For this, the UTC time zone (Coordinated Universal Time, zero offset) is used as the default, taking into account any time zone specified in the string.

Related Topics

- [SQL/JSON Path Expression Item Methods](#)
The Oracle item methods available for a SQL/JSON path expression are described.
- [ISO 8601 Date and Time Support](#)
International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. Oracle Database supports many of the ISO 8601 date and time formats.

14

Clauses Used in SQL Query Functions and Conditions

Clauses `RETURNING`, `wrapper`, `error`, and `empty-field` are described. Each is used in one or more of the SQL functions and conditions `json_value`, `json_query`, `json_table`, `json_serialize`, `json_mergepatch`, `is json`, `is not json`, `json_exists`, and `json_equal`.

- **RETURNING Clause for SQL Query Functions**
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no **RETURNING** clause) are described here.
- **Wrapper Clause for SQL/JSON Query Functions `JSON_QUERY` and `JSON_TABLE`**
SQL/JSON query functions `json_query` and `json_table` accept an optional wrapper clause, which specifies the form of the value returned by `json_query` or used for the data in a `json_table` column. This clause and the default behavior (no wrapper clause) are described here. Examples are provided.
- **Error Clause for SQL Query Functions and Conditions**
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- **Empty-Field Clause for SQL/JSON Query Functions**
SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional **ON EMPTY** clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no **ON EMPTY** clause) are described here.
- **ON MISMATCH Clause for `JSON_VALUE`**
When the **RETURNING** clause specifies a user-defined object-type or collection-type instance, function `json_value` accepts an optional **ON MISMATCH** clause, which specifies handling to use when a targeted JSON value does not match the specified SQL return value. This clause and its default behavior (no **ON MISMATCH** clause) are described here.

14.1 RETURNING Clause for SQL Query Functions

SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no **RETURNING** clause) are described here.

For `json_value`, you can use any of these predefined SQL data types in a **RETURNING** clause: `VARCHAR2`, `NUMBER`, `DATE`, `TIMESTAMP`, `TIMESTAMP WITH TIME ZONE`,

SDO_GEOMETRY, and CLOB. You can also use a user-defined object type or a collection type.

 **Note:**

An instance of Oracle SQL data type `DATE` can include a time component. And in your JSON data you can use a string that represents an ISO 8601 date-with-time value, that is, it can have a time component. But `json_value` with `RETURNING DATE` always returns a SQL `DATE` value that is an *ANSI date*, that is, the value has *no time component*.

To return a SQL `DATE` value that has a time component use `RETURNING TIMESTAMP` and then cast the returned value to `DATE`. For example:

```
SELECT cast(json_value('{ "a" : "2019-01-02T12:34:56" }', '$.a'
RETURNING TIMESTAMP) AS DATE)
FROM DUAL;
```

For `json_query`, `json_serialize`, and `json_mergepatch` you can use `VARCHAR2`, `CLOB`, or `BLOB`. A `BLOB` result is in the AL32UTF8 character set. Whatever the data type returned by `json_serialize`, the returned data represents textual JSON data.

You can optionally specify a length for `VARCHAR2` (default: 4000) and a precision and scale for `NUMBER`.

The default behavior (no `RETURNING` clause) is to use `VARCHAR2(4000)`.

Data type `SDO_GEOMETRY` is for Oracle Spatial and Graph data. In particular, this means that you can use `json_value` with GeoJSON data, which is a format for encoding geographic data in JSON.

The `RETURNING` clause also accepts two optional keywords, `PRETTY` and `ASCII`. If both are present then `PRETTY` must come before `ASCII`. Keyword `PRETTY` is not allowed for `json_value`.

The effect of keyword `PRETTY` is to pretty-print the returned data, by inserting newline characters and indenting. The default behavior is not to pretty-print.

The effect of keyword `ASCII` is to automatically escape all non-ASCII Unicode characters in the returned data, using standard ASCII Unicode escape sequences. The default behavior is not to escape non-ASCII Unicode characters.

 **Tip:**

You can pretty-print the entire context item by using only `$` as the path expression.

If `VARCHAR2` is specified in a `RETURNING` clause then scalars in the value are represented as follows:

- Boolean values are represented by the lowercase strings "true" and "false".

- The `null` value is represented by SQL `NULL`.
- A JSON number is represented in a canonical form. It can thus appear differently in the output string from its representation in textual input data. When represented in canonical form:
 - It can be subject to the precision and range limitations for a SQL `NUMBER`.
 - When it is not subject to the SQL `NUMBER` limitations:
 - * The precision is limited to forty (40) digits.
 - * The optional exponent is limited to nine (9) digits plus a sign (+ or -).
 - * The entire text, including possible signs (-, +), decimal point (.), and exponential indicator (E), is limited to 48 characters.

The **canonical form** of a JSON number:

- Is a JSON number. (It can be parsed in JSON data as a number.)
- Does not have a leading plus (+) sign.
- Has a decimal point (.) only when necessary.
- Has a single zero (0) before the decimal point if the number is a fraction (between zero and one).
- Uses exponential notation (E) only when necessary. In particular, this can be the case if the number of output characters is too limited (by a small *N* for `VARCHAR2(N)`).

Oracle extends the SQL/JSON standard in the case when the returning data type is `VARCHAR2(N)`, by allowing optional keyword **TRUNCATE** immediately after the data type. When **TRUNCATE** is present and the value to return is wider than *N*, the value is truncated — only the first *N* characters are returned. If **TRUNCATE** is absent then this case is treated as an error, handled as usual by an error clause or the default error-handling behavior.

Related Topics

- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- [Using JSON_VALUE To Instantiate a User-Defined Object Type Instance](#)
You can use SQL/JSON function `json_value` to instantiate an instance of a user-defined SQL object type or collection type. You do this by targeting a JSON object or array in the path expression and specifying the object or collection type, respectively, in the `RETURNING` clause.

 **See Also:**

- *Oracle Database SQL Language Reference* for information about SQL data types `DATE` and `TIMESTAMP`
- *Oracle Database SQL Language Reference* for information about SQL data type `NUMBER`
- *Oracle Spatial and Graph Developer's Guide* for information about using Oracle Spatial and Graph data
- GeoJSON.org

14.2 Wrapper Clause for SQL/JSON Query Functions `JSON_QUERY` and `JSON_TABLE`

SQL/JSON query functions `json_query` and `json_table` accept an optional wrapper clause, which specifies the form of the value returned by `json_query` or used for the data in a `json_table` column. This clause and the default behavior (no wrapper clause) are described here. Examples are provided.

The wrapper clause takes one of these forms:

- **WITH WRAPPER** – Use a string value that represents a JSON array containing *all* of the JSON values that match the path expression. The order of the array elements is unspecified.
- **WITHOUT WRAPPER** – Use a string value that represents the *single* JSON *object* or *array* that matches the path expression. Raise an error if the path expression matches either a scalar value (not an object or array) or more than one value.
- **WITH CONDITIONAL WRAPPER** – Use a string value that represents *all* of the JSON values that match the path expression. For zero values, a single scalar value, or multiple values, **WITH CONDITIONAL WRAPPER** is the same as **WITH WRAPPER**. For a single JSON object or array value, it is the same as **WITHOUT WRAPPER**.

The default behavior is **WITHOUT WRAPPER**.

You can add the optional keyword **UNCONDITIONAL** immediately after keyword **WITH**, if you find it clearer: **WITH WRAPPER** and **WITH UNCONDITIONAL WRAPPER** mean the same thing.

You can add the optional keyword **ARRAY** immediately before keyword **WRAPPER**, if you find it clearer: **WRAPPER** and **ARRAY WRAPPER** mean the same thing.

[Table 14-1](#) illustrates the wrapper clause possibilities. The array wrapper is shown in **bold**.

Table 14-1 JSON_QUERY Wrapper Clause Examples

JSON Values Matching Path Expression	WITH WRAPPER	WITHOUT WRAPPER	WITH CONDITIONAL WRAPPER
<code>{"id": 38327}</code> (single object)	<code>[{"id": 38327}]</code>	<code>{"id": 38327}</code>	<code>{"id": 38327}</code>
<code>[42, "a", true]</code> (single array)	<code>[[42, "a", true]]</code>	<code>[42, "a", true]</code>	<code>[42, "a", true]</code>
<code>42</code>	<code>[42]</code>	Error (scalar)	<code>[42]</code>
<code>42, "a", true</code>	<code>[42, "a", true]</code>	Error (multiple values)	<code>[42, "a", true]</code>
none	<code>[]</code>	Error (no values)	<code>[]</code>

Consider, for example, a `json_query` query to retrieve a JSON object. What happens if the path expression matches a JSON scalar value instead of an object, or it matches multiple JSON values (of any kind)? You might want to retrieve the matched values instead of raising an error. For example, you might want to pick one of the values that is an object, for further processing. Using an array wrapper lets you do this.

A conditional wrapper can be convenient if the only reason you are using a wrapper is to avoid raising an error and you do not need to distinguish those error cases from non-error cases. If your application is looking for a single object or array and the data matched by a path expression is just that, then there is no need to wrap that expected value in a singleton array.

On the other hand, with an unconditional wrapper you know that the resulting array is always a wrapper — your application can count on that. If you use a conditional wrapper then your application might need extra processing to interpret a returned array. In Table 14-1, for instance, note that the same array (`[42, "a", true]`) is returned for the very different cases of a path expression matching that array and a path expression matching each of its elements.

14.3 Error Clause for SQL Query Functions and Conditions

Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.

By default, SQL functions and conditions for JSON avoid raising runtime errors. For example, when JSON data is syntactically invalid, `json_exists` and `json_equal` return false and `json_value` returns NULL.

But in some cases you can also specify an error clause, which overrides the default behavior. The error handling you can specify varies, but each SQL function and condition for JSON that lets you specify error handling supports at least the `ERROR ON ERROR` behavior of raising an error.

The optional error clause can take these forms:

- **ERROR ON ERROR** – Raise the error (no special handling).
- **NULL ON ERROR** – Return NULL instead of raising the error.

Not available for `json_exists`.

- **FALSE ON ERROR** – Return false instead of raising the error.
Available *only* for `json_exists` and `json_equal`, for which it is the *default*.
- **TRUE ON ERROR** – Return true instead of raising the error.
Available *only* for `json_exists` and `json_equal`.
- **EMPTY OBJECT ON ERROR** – Return an empty object (`{}`) instead of raising the error.
Available *only* for `json_query`.
- **EMPTY ARRAY ON ERROR** – Return an empty array (`[]`) instead of raising the error.
Available *only* for `json_query`.
- **EMPTY ON ERROR** – Same as `EMPTY ARRAY ON ERROR`.
- **DEFAULT 'literal_return_value' ON ERROR** – Return the specified value instead of raising the error. The value must be a constant at query compile time.

Not available:

- For `json_exists`, `json_equal`, `json_serialize`, `json_mergepatch`, or a `json_table` column value clause that has `json_exists` behavior
- For `json_query` or a `json_table` column value clause that has `json_query` behavior
- For row-level error-handling for `json_table`
- When `SDO_GEOMETRY` is specified either as the `RETURNING` clause data type for `json_value` or as a `json_table` *column* data type

The *default* behavior is `NULL ON ERROR`, except for conditions `json_exists` and `json_equal`.



Note:

There are two levels of error handling for `json_table`, corresponding to its two levels of path expressions: row and column. When present, a column error handler overrides row-level error handling. The default error handler for both levels is `NULL ON ERROR`.



Note:

An `ON EMPTY` clause overrides the behavior specified by `ON ERROR` for the error of trying to match a missing field.

 **Note:**

The `ON ERROR` clause takes effect only for runtime errors that arise when a syntactically correct SQL/JSON path expression is matched against JSON data. A path expression that is syntactically incorrect results in a compile-time syntax error; it is not handled by the `ON ERROR` clause.

Related Topics

- [Empty-Field Clause for SQL/JSON Query Functions](#)
SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional `ON EMPTY` clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no `ON EMPTY` clause) are described here.
- [SQL/JSON Function JSON_TABLE](#)
SQL/JSON function `json_table` projects specific JSON data to columns of various SQL data types. You use it to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.
- [SQL/JSON Function JSON_QUERY](#)
SQL/JSON function `json_query` selects one or more values from JSON data and returns a string (`VARCHAR2`, `CLOB`, or `BLOB` instance) that represents the JSON values. You can thus use `json_query` to retrieve *fragments* of a JSON document.
- [SQL/JSON Function JSON_VALUE](#)
SQL/JSON function `json_value` selects JSON data and returns a SQL scalar or an instance of a user-defined SQL object type or SQL collection type (`varray`, nested table).
- [Oracle SQL Function JSON_SERIALIZE](#)
Oracle SQL function `json_serialize` takes JSON data (of any SQL data type, `VARCHAR2`, `CLOB`, or `BLOB`) as input and returns a textual representation of it (as `VARCHAR2`, `CLOB`, or `BLOB` data). `VARCHAR2(4000)` is the default return type.
- [SQL/JSON Condition JSON_EXISTS](#)
SQL/JSON condition `json_exists` lets you use a SQL/JSON path expression as a row filter, to select rows based on the content of JSON documents. You can use condition `json_exists` in a `CASE` expression or the `WHERE` clause of a `SELECT` statement.
- [ON MISMATCH Clause for JSON_VALUE](#)
When the `RETURNING` clause specifies a user-defined object-type or collection-type instance, function `json_value` accepts an optional `ON MISMATCH` clause, which specifies handling to use when a targeted JSON value does not match the specified SQL return value. This clause and its default behavior (no `ON MISMATCH` clause) are described here.

 **See Also:**

- *Oracle Database SQL Language Reference* for detailed information about the error clause for SQL functions for JSON
- *Oracle Database SQL Language Reference* for detailed information about the error clause for SQL conditions for JSON

14.4 Empty-Field Clause for SQL/JSON Query Functions

SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional `ON EMPTY` clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no `ON EMPTY` clause) are described here.

You generally handle errors for SQL/JSON functions and conditions using an error clause (`ON ERROR`). However, there is a special case where you might want different handling from this general error handling: when querying to match given JSON fields that are missing from the data. Sometimes you do not want to raise an error just because a field to be matched is absent. (A missing field is normally treated as an error.)

You typically use a `NULL ON EMPTY` clause in conjunction with an accompanying `ON ERROR` clause. This combination specifies that other errors are handled according to the `ON ERROR` clause, but the error of trying to match a missing field is handled by just returning `NULL`. If no `ON EMPTY` clause is present then an `ON ERROR` clause handles also the missing-field case.

In addition to `NULL ON EMPTY` there are `ERROR ON EMPTY` and `DEFAULT ... ON EMPTY`, which are analogous to the similarly named `ON ERROR` clauses.

If only an `ON EMPTY` clause is present (no `ON ERROR` clause) then missing-field behavior is specified by the `ON EMPTY` clause, and other errors are handled the same as if `NULL ON ERROR` were present (it is the `ON ERROR` default). If both clauses are absent then only `NULL ON ERROR` is used.

Use `NULL ON EMPTY` for an Index Created on `JSON_VALUE`

`NULL ON EMPTY` is especially useful for the case of a functional index created on a `json_value` expression. The clause has no effect on whether or when the index is picked up, but it is effective in allowing some data to be indexed that would otherwise not be because it is missing a field targeted by the `json_value` expression.

You generally want to use `ERROR ON ERROR` for the queries that populate the index, so that a query path expression that results in multiple values or complex values raises an error. But you sometimes do not want to raise an error just because the field targeted by a path expression is missing — you want that data to be indexed. [Example 26-5](#) illustrates this use of `NULL ON EMPTY` when creating an index on a `json_value` expression.

Related Topics

- [Creating JSON_VALUE Function-Based Indexes](#)
You can create a function-based index for SQL/JSON function `json_value`. You can use the standard syntax for this, explicitly specifying function `json_value`, or you can use the simple dot-notation syntax. Indexes created in either of these ways can be used with both dot-notation queries and `json_value` queries.
- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.

14.5 ON MISMATCH Clause for JSON_VALUE

When the `RETURNING` clause specifies a user-defined object-type or collection-type instance, function `json_value` accepts an optional `ON MISMATCH` clause, which specifies handling to use when a targeted JSON value does not match the specified SQL return value. This clause and its default behavior (no `ON MISMATCH` clause) are described here.

 **Note:**

Clauses `ON ERROR` and `ON EMPTY` apply only when no JSON data matches the path expression. Clause `ON MISMATCH` applies when the JSON data that matches the path expression does not match the specified object or collection return type. Clause `ON MISMATCH` is allowed only when the return type is an instance of a user-defined object type or collection type. If you use it with another return type then a query compile-time error is raised.

When you return an instance of a SQL object or collection type that reflects the JSON data targeted by function `json_value`, the definitions of that targeted data and the object or collection to be returned must match, or else a query compile-time error applies.

The *default* handling of such an error is just to *ignore* it. But you can instead handle such an error in various ways, by providing one or more `ON MISMATCH` clauses, as follows:

- **IGNORE ON MISMATCH** — Explicitly specify the default behavior: ignore the mismatch. The object or collection instance returned can contain one or more SQL `NULL` values because of mismatches against the targeted JSON data.
- **NULL ON MISMATCH** — Return SQL `NULL` as the object-type or collection-type value.
- **ERROR ON MISMATCH** — Raise a query compile-time error for the mismatch.

Each of these `ON MISMATCH` clause types can also be followed, in parentheses `((...))`, by one or more clauses that each indicates a kind of mismatch to handle, separated by commas `(,)`. These are the possible mismatch kinds:

- **MISSING DATA** — Some JSON data was needed to match the object-type or collection-type data, but it was missing.

- **EXTRA DATA** — One or more JSON fields have no corresponding object-type or collection-type data. For example, for JSON field `address` there is no object-type attribute with the same name (matching case-insensitively, by default).
- **TYPE ERROR** — A JSON scalar value has a data type that is incompatible with the corresponding return SQL scalar data type. This can be because of type incompatibility, as put forth in [Table 16-1](#), or because the SQL data type is too constraining (e.g., `VARCHAR(2)` is too short for JSON string `"hello"`).

If no such kind-of-mismatch clause (e.g. `EXTRA DATA`) is present for a given handler (e.g. `NULL ON MISMATCH`) then that handler applies to all kinds of mismatch.

You can have any number of `ON MISMATCH` clauses of different kinds, but if two or more such contradict each other then a query compile-time error is raised.

Related Topics

- [Using JSON_VALUE To Instantiate a User-Defined Object Type Instance](#)
You can use SQL/JSON function `json_value` to instantiate an instance of a user-defined SQL object type or collection type. You do this by targeting a JSON object or array in the path expression and specifying the object or collection type, respectively, in the `RETURNING` clause.

SQL/JSON Condition JSON_EXISTS

SQL/JSON condition `json_exists` lets you use a SQL/JSON path expression as a row filter, to select rows based on the content of JSON documents. You can use condition `json_exists` in a CASE expression or the WHERE clause of a SELECT statement.

Condition `json_exists` checks for the existence of a particular value within JSON data: it returns true if the value is present and false if it is absent. More precisely, `json_exists` returns true if the data it targets matches one or more JSON values. If no JSON values are matched then it returns false.

You can also use `json_exists` to create bitmap indexes for use with JSON data — see [Example 26-1](#).

Error handlers `ERROR ON ERROR`, `FALSE ON ERROR`, and `TRUE ON ERROR` apply. The default is `FALSE ON ERROR`. The handler takes effect when any error occurs, but typically an error occurs when the given JSON data is not well-formed (using lax syntax). Unlike the case for conditions `is json` and `is not json`, condition `json_exists` *expects* the data it examines to be well-formed JSON data.

The second argument to `json_exists` is a SQL/JSON path expression followed by an optional `PASSING` clause and an optional error clause.

The optional filter expression of a SQL/JSON path expression used with `json_exists` can refer to SQL/JSON variables, whose values are passed from SQL by binding them with the `PASSING` clause. The following SQL data types are supported for such variables: `VARCHAR2`, `NUMBER`, `BINARY_DOUBLE`, `DATE`, `TIMESTAMP`, and `TIMESTAMP WITH TIMEZONE`.

 **Note:**

SQL/JSON condition `json_exists` applied to JSON value `null` returns the SQL string `'true'`.

- [Using Filters with JSON_EXISTS](#)
You can use SQL/JSON condition `json_exists` with a path expression that has one or more filter expressions, to select documents that contain matching data. Filters let you test for the existence of documents that have particular fields that satisfy various conditions.
- [JSON_EXISTS as JSON_TABLE](#)
SQL/JSON condition `json_exists` can be viewed as a special case of SQL/JSON function `json_table`.

Related Topics

- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value

returned by the function. This clause and the default behavior (no RETURNING clause) are described here.

- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.



See Also:

Oracle Database SQL Language Reference for information about `json_exists`

15.1 Using Filters with JSON_EXISTS

You can use SQL/JSON condition `json_exists` with a path expression that has one or more filter expressions, to select documents that contain matching data. Filters let you test for the existence of documents that have particular fields that satisfy various conditions.

SQL/JSON condition `json_exists` returns true for documents containing data that matches a SQL/JSON path expression. If the path expression contains a filter, then the data that matches the path to which that filter is applied must also satisfy the filter, in order for `json_exists` to return true for the document containing the data.

A filter applies to the path that immediately precedes it, and the test is whether both (a) the given document has some data that matches that path, and (b) that matching data satisfies the filter. If both of these conditions hold then `json_exists` returns true for the document.

The path expression immediately preceding a filter defines the scope of the patterns used in it. An at-sign (@) within a filter refers to the data targeted by that path, which is termed the *current item* for the filter. For example, in the path expression `$.LineItems?(@.Part.UPCCode == 85391628927)`, @ refers to an occurrence of array `LineItems`.

Example 15-1 JSON_EXISTS: Path Expression Without Filter

This example selects purchase-order documents that have a line item whose part description contains a UPC code entry.

```
SELECT po.po_document FROM j_purchaseorder po
       WHERE json_exists(po.po_document, '$.LineItems.Part.UPCCode');
```

Example 15-2 JSON_EXISTS: Current Item and Scope in Path Expression Filters

This example shows three *equivalent* ways to select documents that have a line item whose part contains a UPC code with a value of 85391628927.

```
SELECT po.po_document FROM j_purchaseorder po
       WHERE json_exists(po.po_document, '$?(@.LineItems.Part.UPCCode ==
7 85391628927)');
```

```
SELECT po.po_document FROM j_purchaseorder po
WHERE json_exists(po.po_document, '$.LineItems?(@.Part.UPCCode ==
85391628927)');
```

```
SELECT po.po_document FROM j_purchaseorder po
WHERE json_exists(po.po_document, '$.LineItems.Part?(@.UPCCode ==
85391628927)');
```

- In the first query, the scope of the filter is the context item, that is, an entire purchase order. @ refers to the context item.
- In the second query, the filter scope is a `LineItems` array (and each of its elements, implicitly). @ refers to an element of that array.
- In the third query, the filter scope is a `Part` field of an element in a `LineItems` array. @ refers to a `Part` field.

Example 15-3 JSON_EXISTS: Filter Conditions Depend On the Current Item

This example selects purchase-order documents that have both a line item with a part that has UPC code 85391628927 *and* a line item with an order quantity greater than 3. The scope of each filter, that is, the current item, is in this case the context item. Each filter condition applies independently (to the same document); the two conditions do *not* necessarily apply to the *same* line item.

```
SELECT po.po_document FROM j_purchaseorder po
WHERE json_exists(po.po_document, '$?(@.LineItems.Part.UPCCode ==
85391628927
&& @.LineItems.Quantity > 3)');
```

Example 15-4 JSON_EXISTS: Filter Downscoping

This example looks similar to [Example 15-3](#), but it acts quite differently. It selects purchase-order documents that have a line item with a part that has UPC code *and with* an order quantity greater than 3. The scope of the current item in the filter is at a lower level; it is not the context item but a `LineItems` array element. That is, the *same line item* must satisfy both conditions, for `json_exists` to return true.

```
SELECT po.po_document FROM j_purchaseorder po
WHERE json_exists(po.po_document, '$.LineItems?(@.Part.UPCCode ==
85391628927
&& @.Quantity > 3)');
```

Example 15-5 JSON_EXISTS: Path Expression Using Path-Expression exists Condition

This example shows how to downscope one part of a filter while leaving another part scoped at the document (context-item) level. It selects purchase-order documents that have a `User` field whose value is "ABULL" and documents that have a line item with a part that has UPC code *and with* an order quantity greater than 3. That is, it selects the same documents selected by [Example 15-4](#), as well as all documents that have "ABULL" as the user. The argument to path-expression predicate `exists` is a path expression that specifies particular line items; the predicate returns true if a match is found, that is, if any such line items exist.

(If you use this example or similar with SQL*Plus then you must use SET DEFINE OFF first, so that SQL*Plus does not interpret && exists as a substitution variable and prompt you to define it.)

```
SELECT po.po_document FROM j_purchaseorder po
WHERE json_exists(po.po_document,
                  '$?(@.User == "ABULL"
                  && exists(@.LineItems?(@.Part.UPCCode ==
85391628927
                  && @.Quantity > 3)))');
```

Related Topics

- [Basic SQL/JSON Path Expression Syntax](#)
The basic syntax of a SQL/JSON path expression is presented. It is composed of a context-item symbol (\$) followed by zero or more object, array, and descendant steps, each of which can be followed by a filter expression, followed optionally by a function step. Examples are provided.

15.2 JSON_EXISTS as JSON_TABLE

SQL/JSON condition `json_exists` can be viewed as a special case of SQL/JSON function `json_table`.

Example 15-6 illustrates the equivalence: the two `SELECT` statements have the same effect.

In addition to perhaps helping you understand `json_exists` better, this equivalence is important practically, because it means that you can use either to get the same effect.

In particular, if you use `json_exists` more than once, or you use it in combination with `json_value` or `json_query` (which can also be expressed using `json_table`), to access the same data, then a single invocation of `json_table` presents the advantage that the data is parsed only once.

Because of this, the optimizer often automatically rewrites multiple invocations of `json_exists`, `json_value` and `json_query` (any combination) to fewer invocations of `json_table`.

Example 15-6 JSON_EXISTS Expressed Using JSON_TABLE

```
SELECT select_list
FROM table WHERE json_exists(column, json_path error_handler ON
ERROR);

SELECT select_list
FROM table,
      json_table(column, '$' error_handler ON ERROR
                COLUMNS ("COLUMN_ALIAS" NUMBER EXISTS PATH json_path)) AS "JT"
WHERE jt.column_alias = 1;
```


Related Topics

- [JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions](#)
SQL/JSON function `json_table` generalizes SQL/JSON condition `json_exists` and SQL/JSON functions `json_value` and `json_query`. Everything that you can do using these functions you can do using `json_table`. For the jobs they accomplish, the syntax of these functions is simpler to use than is the syntax of `json_table`.

SQL/JSON Function JSON_VALUE

SQL/JSON function `json_value` selects JSON data and returns a SQL scalar or an instance of a user-defined SQL object type or SQL collection type (varray, nested table).

- If `json_value` targets a single *scalar* JSON value then it returns a scalar SQL value. You can specify the SQL data type for the returned scalar value. By default it is `VARCHAR2(4000)`.
- If `json_value` targets a JSON *array*, and you specify a SQL *collection type* (varray or nested table) as the return type, then `json_value` returns an instance of that collection type.

The elements of a targeted JSON array provide the elements of the returned collection-type instance. A scalar JSON array element produces a scalar SQL value in the returned collection instance (see previous). A JSON array element that is an object (see next) or an array is handled recursively.

- If `json_value` targets a JSON *object*, and you specify a user-defined SQL *object type* as the return type, then `json_value` returns an instance of that object type.

The field values of a targeted JSON object provide the attribute values of the returned object-type instance. The field names of the targeted JSON object are compared with the SQL names of the SQL object attributes. A scalar field value produces a scalar SQL value in the returned object-type instance (see above). A field value that is an array (see previous) or an object is handled recursively,

Ultimately it is the names of JSON fields with scalar values that are compared with the names of scalar SQL object attributes. If the names do not match exactly, case-sensitively, then a *mismatch error* is raised at query compile time.

You can also use `json_value` to create function-based B-tree indexes for use with JSON data — see [Indexes for JSON Data](#).

Function `json_value` has two required arguments, and it accepts optional returning and error clauses.

The first argument to `json_value` is a SQL expression that returns an instance of either a scalar SQL data type or a user-defined SQL object type. A scalar return value can be of data type `VARCHAR2`, `BLOB`, or `CLOB`.

The first argument can be a table or view column value, a PL/SQL variable, or a bind variable with proper casting. The result of evaluating the SQL expression is used as the *context item* for evaluating the path expression.

The second argument to `json_value` is a SQL/JSON path expression followed by optional clauses `RETURNING`, `ON ERROR`, and `ON EMPTY`. The path expression must target a single scalar value, or else an error occurs.

The *default* error-handling behavior is `NULL ON ERROR`, which means that no value is returned if an error occurs — an error is not raised. In particular, if the path expression targets a non-scalar value, such as an array, no error is raised, by default. To ensure that an error is raised, use `ERROR ON ERROR`.

 **Note:**

Each field name in a given JSON object is not necessarily unique; the same field name may be repeated. The streaming evaluation that Oracle Database employs always uses only one of the object members that have a given field name; any other members with the same field name are ignored. It is unspecified which of multiple such members is used.

See also [Unique Versus Duplicate Fields in JSON Objects](#).

- [Using SQL/JSON Function `JSON_VALUE` With a Boolean JSON Value](#)
JSON has the Boolean values `true` and `false`. When SQL/JSON function `json_value` evaluates a SQL/JSON path expression and the result is JSON `true` or `false`, it can be returned to PL/SQL as a `BOOLEAN` value, or it can be returned to SQL as the `VARCHAR2` value `'true'` or `'false'`.
- [SQL/JSON Function `JSON_VALUE` Applied to a null JSON Value](#)
SQL/JSON function `json_value` applied to JSON value `null` returns SQL `NULL`, not the SQL string `'null'`. This means, in particular, that you cannot use `json_value` to distinguish the JSON value `null` from the absence of a value; SQL `NULL` indicates both cases.
- [Using `JSON_VALUE` To Instantiate a User-Defined Object Type Instance](#)
You can use SQL/JSON function `json_value` to instantiate an instance of a user-defined SQL object type or collection type. You do this by targeting a JSON object or array in the path expression and specifying the object or collection type, respectively, in the `RETURNING` clause.
- [`JSON_VALUE` as `JSON_TABLE`](#)
SQL/JSON function `json_value` can be viewed as a special case of function `json_table`.

Related Topics

- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional `RETURNING` clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no `RETURNING` clause) are described here.
- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- [Empty-Field Clause for SQL/JSON Query Functions](#)
SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional `ON EMPTY` clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no `ON EMPTY` clause) are described here.

 **See Also:**

Oracle Database SQL Language Reference for information about `json_value`

16.1 Using SQL/JSON Function JSON_VALUE With a Boolean JSON Value

JSON has the Boolean values `true` and `false`. When SQL/JSON function `json_value` evaluates a SQL/JSON path expression and the result is JSON `true` or `false`, it can be returned to PL/SQL as a `BOOLEAN` value, or it can be returned to SQL as the `VARCHAR2` value `'true'` or `'false'`.

In PL/SQL code, `BOOLEAN` is a valid PL/SQL return type for built-in PL/SQL function `json_value`. [Example 16-1](#) illustrates this.

Oracle SQL has no Boolean data type, so a string (`VARCHAR2`) value is used to return a JSON Boolean value. [Example 16-2](#) illustrates this — the query returns the string `'true'`.

SQL/JSON function `json_table` generalizes other SQL/JSON query functions such as `json_value`. When you use it to project a JSON Boolean value, `json_value` is used implicitly, and the resulting SQL value is returned as a `VARCHAR2` value. The data type of the projection column must therefore be `VARCHAR2`.

Example 16-1 JSON_VALUE: Returning a JSON Boolean Value to PL/SQL as BOOLEAN

PL/SQL also has exception handling. This example uses clause `ERROR ON ERROR`, to raise an error (which can be handled by user code) in case of error.

```
DECLARE
  b BOOLEAN;
  jsonData CLOB;
BEGIN
  SELECT po_document INTO jsonData FROM j_purchaseorder WHERE rownum =
  1;
  b := json_value(jsonData, '$.AllowPartialShipment'
                 RETURNING BOOLEAN
                 ERROR ON ERROR);
END;
/
```

Example 16-2 JSON_VALUE: Returning a JSON Boolean Value to SQL as VARCHAR2

```
SELECT json_value(po_document, '$.AllowPartialShipment')
       FROM j_purchaseorder;
```

Related Topics

- [JSON_VALUE as JSON_TABLE](#)
SQL/JSON function `json_value` can be viewed as a special case of function `json_table`.
- [JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions](#)
SQL/JSON function `json_table` generalizes SQL/JSON condition `json_exists` and SQL/JSON functions `json_value` and `json_query`. Everything that you can do using these functions you can do using `json_table`. For the jobs they accomplish, the syntax of these functions is simpler to use than is the syntax of `json_table`.

16.2 SQL/JSON Function `JSON_VALUE` Applied to a null JSON Value

SQL/JSON function `json_value` applied to JSON value `null` returns SQL `NULL`, not the SQL string `'null'`. This means, in particular, that you cannot use `json_value` to distinguish the JSON value `null` from the absence of a value; SQL `NULL` indicates both cases.

16.3 Using `JSON_VALUE` To Instantiate a User-Defined Object Type Instance

You can use SQL/JSON function `json_value` to instantiate an instance of a user-defined SQL object type or collection type. You do this by targeting a JSON object or array in the path expression and specifying the object or collection type, respectively, in the `RETURNING` clause.

The elements of a targeted JSON array provide the elements of a returned collection-type instance. The JSON array elements must correspond, one-to-one, with the collection-type elements. If they do not then a mismatch error occurs. A JSON array element that is an object (see next) or an array is handled recursively.

The fields of a targeted JSON object provide the attribute values of a returned object-type instance. The JSON fields must correspond, one-to-one, with the object-type attributes. If they do not then a mismatch error occurs.

The field names of the targeted JSON object are compared with the SQL names of the object attributes. A field value that is an array or an object is handled recursively, so that ultimately it is the names of JSON fields with scalar values that are compared with the names of scalar SQL object attributes. If the names do not match (case insensitively, by default), then a mismatch error occurs.

If all names match then the corresponding data types are checked for compatibility. If there is any type incompatibility then a mismatch error occurs. [Table 16-1](#) specifies the compatible scalar data types — any other type combinations are incompatible entails a mismatch error.

Table 16-1 Compatible Scalar Data Types: Converting JSON to SQL

JSON Type (Source)	SQL Type (Destination)	Notes
<code>string</code>	<code>VARCHAR2</code>	None

Table 16-1 (Cont.) Compatible Scalar Data Types: Converting JSON to SQL

JSON Type (Source)	SQL Type (Destination)	Notes
string	CLOB	None
string	NUMBER	The JSON string must be numeric.
string	DATE	The JSON string must have a supported ISO 8601 format.
string	TIMESTAMP	The JSON string must have a supported ISO 8601 format.
number	NUMBER	None
number	VARCHAR2	None
number	CLOB	None
boolean	VARCHAR2	The instance value is the SQL string "true" or "false".
boolean	CLOB	The instance value is the SQL string "true" or "false".
null	Any SQL data type.	The instance value is SQL NULL.

A *mismatch error* occurs at query compile time if any of the following are true. By *default*, mismatch errors are *ignored*, but you can change this error handling by including one or more `ON MISMATCH` clauses in your invocation of `json_value`.

- The fields of a targeted JSON object, or the elements of a targeted JSON array, do not *correspond in number and kind* to the attributes of the specified object-type instance, or to the elements of the specified collection-type instance, respectively.
- The fields of a targeted JSON object do not have the *same names* as the attributes of a specified object-type instance. By default this matching is case-insensitive.
- The JSON and SQL *scalar data types* of a JSON value and its corresponding object attribute value or collection element value are not *compatible*, according to [Table 16-1](#).

Example 16-3 Instantiate a User-Defined Object Instance From JSON Data with JSON_VALUE

This example defines SQL object types `shipping_t` and `addr_t`. Object type `shipping_t` has attributes `name` and `address`, which have types `VARCHAR2(30)` and `addr_t`, respectively.

Object type `addr_t` has attributes `street` and `city`.

The example uses `json_value` to select the JSON object that is the value of field `ShippingInstructions` and return an instance of SQL object type `shipping_t`. Names of the object-type attributes are matched against JSON object field names *case-insensitively*, so that, for example, attribute `address` (which is the same as `ADDRESS`) of SQL object-type `shipping_t` matches JSON field `address`.

(The query output is shown pretty-printed here, for clarity.)

```
CREATE TYPE shipping_t AS OBJECT
  (name    VARCHAR2(30),
   address addr_t);

CREATE TYPE addr_t AS OBJECT
  (street VARCHAR2(100),
   city   VARCHAR2(30));

-- Query data to return shipping_t instances:
SELECT json_value(po_document, '$.ShippingInstructions'
                 RETURNING shipping_t)
       FROM j_purchaseorder;

JSON_VALUE(PO_DOCUMENT, '$.SHIPPINGINSTRUCTIONS' RETURNING
-----
SHIPPING_T('Alexis Bull',
           ADDR_T('200 Sporting Green',
                 'South San Francisco'))
SHIPPING_T('Sarah Bell',
           ADDR_T('200 Sporting Green',
                 'South San Francisco'))
```

Example 16-4 Instantiate a Collection Type Instance From JSON Data with JSON_VALUE

This example defines SQL collection type `items_t` and SQL object types `part_t` and `item_t`. An instance of collection type `items_t` is a varray of `item_t` instances. Attribute `part` of object-type `item_t` is itself of SQL object-type `part_t`.

It then uses `json_value` to select the JSON

(The query output is shown pretty-printed here, for clarity.)

```
CREATE TYPE part_t AS OBJECT
  (description VARCHAR2(30),
   unitprice   NUMBER);

CREATE TYPE item_t AS OBJECT
  (itemnumber NUMBER,
   part       part_t);

CREATE TYPE items_t AS VARRAY(10) OF item_t;

-- Query data to return items_t collections of item_t objects
SELECT json_value(po_document, '$.LineItems' RETURNING items_t)
       FROM j_purchaseorder;

JSON_VALUE(PO_DOCUMENT, '$.LINEITEMS' RETURNING ITEMS_T USIN
-----
ITEMS_T(ITEM_T(1, PART_T('One Magic Christmas', 19.95)),
        ITEM_T(2, PART_T('Lethal Weapon', 19.95)))
ITEMS_T(ITEM_T(1, PART_T('Making the Grade', 20)),
```

```
ITEM_T(2, PART_T('Nixon', 19.95)),
ITEM_T(3, PART_T(NULL, 19.95))
```

Related Topics

- [ON MISMATCH Clause for JSON_VALUE](#)
When the RETURNING clause specifies a user-defined object-type or collection-type instance, function json_value accepts an optional ON MISMATCH clause, which specifies handling to use when a targeted JSON value does not match the specified SQL return value. This clause and its default behavior (no ON MISMATCH clause) are described here.

See Also:

Oracle Database SQL Language Reference for information about json_value

16.4 JSON_VALUE as JSON_TABLE

SQL/JSON function json_value can be viewed as a special case of function json_table.

[Example 16-5](#) illustrates the equivalence: the two SELECT statements have the same effect.

In addition to perhaps helping you understand json_value better, this equivalence is important practically, because it means that you can use either function to get the same effect.

In particular, if you use json_value more than once, or you use it in combination with json_exists or json_query (which can also be expressed using json_table), to access the same data, then a single invocation of json_table presents the advantage that the data is parsed only once.

Because of this, the optimizer often automatically rewrites multiple invocations of json_exists, json_value and json_query (any combination) to fewer invocations of json_table.

Example 16-5 JSON_VALUE Expressed Using JSON_TABLE

```
SELECT json_value(column, json_path RETURNING data_type error_handler ON
ERROR)
FROM table;
```

```
SELECT jt.column_alias
FROM table,
json_table(column, '$' error_handler ON ERROR
COLUMNS ("COLUMN_ALIAS" data_type PATH json_path)) AS "JT";
```


Related Topics

- [JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions](#)
SQL/JSON function `json_table` generalizes SQL/JSON condition `json_exists` and SQL/JSON functions `json_value` and `json_query`. Everything that you can do using these functions you can do using `json_table`. For the jobs they accomplish, the syntax of these functions is simpler to use than is the syntax of `json_table`.

SQL/JSON Function JSON_QUERY

SQL/JSON function `json_query` selects one or more values from JSON data and returns a string (`VARCHAR2`, `CLOB`, or `BLOB` instance) that represents the JSON values. You can thus use `json_query` to retrieve *fragments* of a JSON document.

The first argument to `json_query` is a SQL expression that returns an instance of a scalar SQL data type (that is, not an object or collection data type). It can be of data type `VARCHAR2`, `CLOB`, or `BLOB`. It can be a table or view column value, a PL/SQL variable, or a bind variable with proper casting. The result of evaluating the SQL expression is used as the *context item* for evaluating the path expression.

The second argument to `json_query` is a SQL/JSON path expression followed by optional clauses `RETURNING`, `WRAPPER`, `ON ERROR`, and `ON EMPTY`. The path expression can target any number of JSON values.

In the `RETURNING` clause you can specify data type `VARCHAR2`, `CLOB`, or `BLOB`. A `BLOB` result is in the AL32UTF8 character set. (`VARCHAR2` is the default.) The value returned always contains well-formed JSON data. This includes ensuring that non-ASCII characters in string values are escaped as needed. For example, an ASCII TAB character (Unicode character CHARACTER TABULATION, U+0009) is escaped as `\t`. Keywords `FORMAT JSON` are not needed (or available) for `json_query` — JSON formatting is implicit for the return value.

The wrapper clause determines the form of the returned string value.

The error clause for `json_query` can specify `EMPTY ON ERROR`, which means that an empty array (`[]`) is returned in case of error (no error is raised).

[Example 17-1](#) shows an example of the use of SQL/JSON function `json_query` with an array wrapper. For each document it returns a `VARCHAR2` value whose contents represent a JSON array with elements the phone types, in an unspecified order. For the document in [Example 4-2](#) the phone types are "Office" and "Mobile", and the array returned is either `["Mobile", "Office"]` or `["Office", "Mobile"]`.

Note that if path expression `$.ShippingInstructions.Phone.type` were used in [Example 17-1](#) it would give the same result. Because of SQL/JSON path-expression syntax relaxation, `[*].type` is equivalent to `.type`.

 **See Also:**

Oracle Database SQL Language Reference for information about `json_query`

Example 17-1 Selecting JSON Values Using JSON_QUERY

```
SELECT json_query(po_document, '$.ShippingInstructions.Phone[*].type'  
                WITH WRAPPER)  
FROM j_purchaseorder;
```

- [JSON_QUERY as JSON_TABLE](#)
SQL/JSON function `json_query` can be viewed as a special case of function `json_table`.

Related Topics

- [SQL/JSON Path Expression Syntax Relaxation](#)
The basic SQL/JSON path-expression syntax is relaxed to allow implicit array wrapping and unwrapping. This means that you need not change a path expression in your code if your data evolves to replace a JSON value with an array of such values, or vice versa. Examples are provided.
- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no **RETURNING** clause) are described here.
- [Wrapper Clause for SQL/JSON Query Functions JSON_QUERY and JSON_TABLE](#)
SQL/JSON query functions `json_query` and `json_table` accept an optional wrapper clause, which specifies the form of the value returned by `json_query` or used for the data in a `json_table` column. This clause and the default behavior (no wrapper clause) are described here. Examples are provided.
- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- [Empty-Field Clause for SQL/JSON Query Functions](#)
SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional **ON EMPTY** clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no **ON EMPTY** clause) are described here.

17.1 JSON_QUERY as JSON_TABLE

SQL/JSON function `json_query` can be viewed as a special case of function `json_table`.

Example 17-2 illustrates the equivalence: the two `SELECT` statements have the same effect.

In addition to perhaps helping you understand `json_query` better, this equivalence is important practically, because it means that you can use either function to get the same effect.

In particular, if you use `json_query` more than once, or you use it in combination with `json_exists` or `json_value` (which can also be expressed using `json_table`), to

access the same data, then a single invocation of `json_table` presents the advantage that the data is parsed only once.

Because of this, the optimizer often automatically rewrites multiple invocations of `json_exists`, `json_value` and `json_query` (any combination) to fewer invocations of `json_table`.

Example 17-2 JSON_QUERY Expressed Using JSON_TABLE

```
SELECT json_query(column, json_path
                RETURNING data_type array_wrapper error_handler ON
ERROR)
FROM table;
```

```
SELECT jt.column_alias
FROM table,
     json_table(column, '$' error_handler ON ERROR
               COLUMNS ("COLUMN_ALIAS" data_type FORMAT JSON array_wrapper
                        PATH json_path)) AS "JT";
```

Related Topics

- [JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions](#)
SQL/JSON function `json_table` generalizes SQL/JSON condition `json_exists` and SQL/JSON functions `json_value` and `json_query`. Everything that you can do using these functions you can do using `json_table`. For the jobs they accomplish, the syntax of these functions is simpler to use than is the syntax of `json_table`.

18

SQL/JSON Function `JSON_TABLE`

SQL/JSON function `json_table` projects specific JSON data to columns of various SQL data types. You use it to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.

You can then insert this virtual table into a pre-existing database table, or you can query it using SQL — in a join expression, for example.

A common use of `json_table` is to create a *view* of JSON data. You can use such a view just as you would use any table or view. This lets applications, tools, and programmers operate on JSON data without consideration of the syntax of JSON or JSON path expressions.

Defining a view over JSON data in effect maps a kind of *schema* onto that data. This mapping is *after the fact*: the underlying JSON data can be defined and created without any regard to a schema or any particular pattern of use. Data first, schema later.

Such a schema (mapping) imposes no restriction on the kind of JSON documents that can be stored in the database (other than being well-formed JSON data). The view exposes only data that conforms to the mapping (schema) that defines the view. To change the schema, just redefine the view — no need to reorganize the underlying JSON data.

You use `json_table` in a SQL `FROM` clause. It is a **row source**: it generates a row of virtual-table data for each JSON value selected by a *row path expression* (row pattern). The columns of each generated row are defined by the *column path expressions* of the `COLUMNS` clause.

Typically a `json_table` invocation is laterally joined, implicitly, with a source table in the `FROM` list, whose rows each contain a JSON document that is used as input to the function. `json_table` generates zero or more new rows, as determined by evaluating the row path expression against the input document.

The first argument to `json_table` is a SQL expression. It can be a table or view column value, a PL/SQL variable, or a bind variable with proper casting. The result of evaluating the expression is used as the *context item* for evaluating the row path expression.

The second argument to `json_table` is the SQL/JSON row path expression followed by an optional error clause for handling the row and the (required) `COLUMNS` clause, which defines the columns of the virtual table to be created. There is no `RETURNING` clause.

There are two levels of error handling for `json_table`, corresponding to the two levels of path expressions: row and column. When present, a column error handler overrides row-level error handling. The default error handler for both levels is `NULL ON ERROR`.

As an alternative to passing the context-item argument and the row path expression, you can use simple dot-notation syntax. (You can still use an error clause, and the `COLUMNS` clause is still required.) Dot notation specifies a table or view column together

with a simple path to the targeted JSON data. For example, these two queries are equivalent:

```
json_table(t.j, '$.ShippingInstructions.Phone[*]' ...)
```

```
json_table(t.j.ShippingInstructions.Phone[*] ...)
```

And in cases where the row path expression is only '\$', which targets the entire document, you can omit the path part. These queries are equivalent:

```
json_table(t.j, '$' ...)
```

```
json_table(t.j ...)
```

[Example 18-1](#) illustrates the difference between using the simple dot notation and using the fuller, more explicit notation.

Example 18-1 Equivalent JSON_TABLE Queries: Simple and Full Syntax

This example uses `json_table` for two equivalent queries. The first query uses the simple, dot-notation syntax for the expressions that target the row and column data. The second uses the full syntax.

Except for column `Special Instructions`, whose SQL identifier is quoted, the SQL column names are, in effect, uppercase. (Identifier `Special Instructions` contains a space character.)

In the first query the column names are written exactly the same as the names of the targeted object fields, including with respect to letter case. Regardless of whether they are quoted, they are interpreted case-sensitively for purposes of establishing the default path (the path used when there is no explicit `PATH` clause).

The second query has:

- Separate arguments of a JSON column-expression and a SQL/JSON row path-expression
- Explicit column data types of `VARCHAR2(4000)`
- Explicit `PATH` clauses with SQL/JSON column path expressions, to target the object fields that are projected

```
SELECT jt.*
  FROM j_purchaseorder po,
       json_table(po.po_document
                 COLUMNS ("Special Instructions",
                          NESTED LineItems[*]
                          COLUMNS (ItemNumber NUMBER,
                                   Description PATH Part.Description))) AS
"JT";
```

```
SELECT jt.*
  FROM j_purchaseorder po,
       json_table(po.po_document,
                 '$')
```

```

        COLUMNS (
            "Special Instructions" VARCHAR2(4000) PATH '$.Special
Instructions"',
        NESTED PATH '$.LineItems[*]'
        COLUMNS (
            ItemNumber NUMBER          PATH '$.ItemNumber',
            Description VARCHAR(4000)  PATH '$.Part.Description'))
AS "JT";

```

- [SQL NESTED Clause Instead of JSON_TABLE](#)
In a `SELECT` clause you can often use a `NESTED` clause instead of SQL/JSON function `json_table`. This can mean a simpler query expression. It also has the advantage of including rows with non-NULL relational columns when the JSON column is NULL.
- [COLUMNS Clause of SQL/JSON Function JSON_TABLE](#)
The mandatory `COLUMNS` clause for SQL/JSON function `json_table` defines the columns of the virtual table that the function creates.
- [JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions](#)
SQL/JSON function `json_table` generalizes SQL/JSON condition `json_exists` and SQL/JSON functions `json_value` and `json_query`. Everything that you can do using these functions you can do using `json_table`. For the jobs they accomplish, the syntax of these functions is simpler to use than is the syntax of `json_table`.
- [Using JSON_TABLE with JSON Arrays](#)
A JSON value can be an array or can include one or more arrays, nested to any number of levels inside other JSON arrays or objects. You can use a `json_table` `NESTED` path clause to project specific elements of an array.
- [Creating a View Over JSON Data Using JSON_TABLE](#)
To improve query performance you can create a view over JSON data that you project to columns using SQL/JSON function `json_table`. To further improve query performance you can create a *materialized view* and place the JSON data *in memory*.

Related Topics

- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- [SQL/JSON Function JSON_QUERY](#)
SQL/JSON function `json_query` selects one or more values from JSON data and returns a string (VARCHAR2, CLOB, or BLOB instance) that represents the JSON values. You can thus use `json_query` to retrieve *fragments* of a JSON document.

See Also:

Oracle Database SQL Language Reference for information about `json_table`

18.1 SQL NESTED Clause Instead of JSON_TABLE

In a `SELECT` clause you can often use a `NESTED` clause instead of SQL/JSON function `json_table`. This can mean a simpler query expression. It also has the advantage of including rows with non-`NULL` relational columns when the JSON column is `NULL`.

The `NESTED` clause is a shortcut for using `json_table` with an ANSI left outer join. That is, these two queries are equivalent:

```
SELECT ...
  FROM mytable NESTED jcol COLUMNS (...);
```

```
SELECT ...
  FROM mytable t1 LEFT OUTER JOIN
    json_table(t1.jcol COLUMNS (...))
  ON 1=1;
```

Using a left outer join with `json_table`, or using the `NESTED` clause, allows the selection result to include rows with relational columns where there is no corresponding JSON-column data, that is, where the JSON column is `NULL`. The only semantic difference between the two is that if you use a `NESTED` clause then the JSON column itself is not included in the result.

The `NESTED` clause provides the same `COLUMNS` clause as `json_table`, including the possibility of nested columns. Advantages of using `NESTED` are that you need not provide an `is json` check constraint (needed for `json_table` with the simple dot notation) or a table alias, and you need not specify `LEFT OUTER JOIN`. The `NESTED` clause syntax is simpler, it allows all of the flexibility of the `COLUMNS` clause, and it performs an implicit left outer join.

Example 18-2 Equivalent: SQL NESTED and JSON_TABLE with LEFT OUTER JOIN

These two queries are equivalent. One uses SQL/JSON function `json_table` with an explicit `LEFT OUTER JOIN`. The other uses a SQL `NESTED` clause.

```
SELECT id, requestor, type, "number"
  FROM j_purchaseorder LEFT OUTER JOIN
    json_table(po_document
      COLUMNS (Requestor,
        NESTED ShippingInstructions.Phone[*]
          COLUMNS (type, "number")))
  ON 1=1);
```

```
SELECT id, requestor, type, "number"
  FROM j_purchaseorder NESTED
    po_document
      COLUMNS (Requestor,
        NESTED ShippingInstructions.Phone[*]
          COLUMNS (type, "number"));
```


The output is the same in both cases:

```
7C3A54B183056369E0536DE05A0A15E4 Alexis Bull Office 909-555-7307
7C3A54B183056369E0536DE05A0A15E4 Alexis Bull Mobile 415-555-1234
7C3A54B183066369E0536DE05A0A15E4 Sarah Bell
```

If table `j_purchaseorder` had a row with non-NULL values for columns `id` and `requestor`, but a NULL value for column `po_document` then that row would appear in both cases. But it would not appear in the `json_table` case if `LEFT OUTER JOIN` were absent.

18.2 COLUMNS Clause of SQL/JSON Function JSON_TABLE

The mandatory `COLUMNS` clause for SQL/JSON function `json_table` defines the columns of the virtual table that the function creates.

It consists of the keyword `COLUMNS` followed by the following entries, enclosed in parentheses. Other than the optional `FOR ORDINALITY` entry, each entry in the `COLUMNS` clause is either a *regular* column specification or a *nested* columns specification.

- At most one entry in the `COLUMNS` clause can be a column name followed by the keywords `FOR ORDINALITY`, which specifies a column of generated row numbers (SQL data type `NUMBER`). These numbers start with one. For example:

```
COLUMNS (linenum FOR ORDINALITY, ProductID)
```

- A **regular column** specification consists of a column name followed by an optional scalar data type for the column, which can be SQL data type `VARCHAR2`, `NUMBER`, `DATE`, `TIMESTAMP`, `TIMESTAMP WITH TIME ZONE`, or `SDO_GEOMETRY` (the same as for the `RETURNING` clause of `json_value`), followed by an optional value clause and an optional `PATH` clause. The default data type is `VARCHAR2(4000)`.

Data type `SDO_GEOMETRY` is used for Oracle Spatial and Graph data. In particular, this means that you can use `json_table` with GeoJSON data, which is a format for encoding geographic data in JSON.

Oracle extends the SQL/JSON standard in the case when the returning data type for a column is `VARCHAR2(N)`, by allowing optional keyword `TRUNCATE` immediately after the data type. When `TRUNCATE` is present and the value to return is wider than `N`, the value is truncated — only the first `N` characters are returned. If `TRUNCATE` is absent then this case is treated as an error, handled as usual by an error clause or the default error-handling behavior.

- A **nested columns** specification consists of the keyword `NESTED` followed by an optional `PATH` keyword, a SQL/JSON row path expression, and then a `COLUMNS` clause. This `COLUMNS` clause specifies columns that represent nested data. The row path expression used here provides a refined context for the specified nested columns: each nested column path expression is relative to the row path expression. You can nest columns clauses to project values that are present in arrays at different levels to columns of the same row.

A `COLUMNS` clause at any level (nested or not) has the same characteristics. In other words, the `COLUMNS` clause is defined recursively. For each level of nesting

(that is, for each use of keyword `NESTED`), the nested `COLUMNS` clause is said to be the **child** of the `COLUMNS` clause within which it is nested, which is its **parent**. Two or more `COLUMNS` clauses that have the same parent clause are **siblings**.

The virtual tables defined by parent and child `COLUMNS` clauses are joined using an *outer* join, with the parent being the outer table. The virtual columns defined by sibling `COLUMNS` clauses are joined using a *union* join.

[Example 18-1](#) and [Example 18-8](#) illustrate the use of a nested columns clause.

The only thing required in a regular column specification is the column name. Defining the column projection in more detail, by specifying a scalar data type, value handling, or a target path, is optional.

- The optional **value** clause specifies how to handle the data projected to the column: whether to handle it as would `json_value`, `json_exists`, or `json_query`. This value handling includes the return data type, return format (`pretty` or `ASCII`), wrapper, and error treatment.

By default, the projected data is handled as if by `json_value`. If you use keyword **EXISTS** then it is handled as if by `json_exists`. If you use keywords **FORMAT JSON** then it is handled as if by `json_query`.

For `FORMAT JSON` you can override the default wrapping behavior by adding an explicit wrapper clause.

You can override the default error handling for the given handler (`json_value`, `json_exists`, or `json_query`) by adding an explicit error clause appropriate for it.

- The optional **PATH** clause specifies the portion of the row that is to be used as the column content. The column path expression following keyword `PATH` is matched against the context item provided by the virtual row. The column path expression must represent a *relative* path; it is relative to the path specified by the row path expression.

If the `PATH` clause is not present then the behavior is the same as if it were present with a path of `'$.<column-name>'`, where `<column-name>` is the column name.

That is, the name of the object field that is targeted is taken implicitly as the column name. For purposes of specifying the targeted field *only*, the SQL identifier used for `<column-name>` is interpreted *case-sensitively*, even if it is not quoted. The SQL name of the column follows the usual rule: if it is enclosed in double quotation marks (") then the letter case used is significant; otherwise, it is not (it is treated as if uppercase).

For example, these two `COLUMNS` clauses are equivalent. For SQL, case is significant *only* for column `Comments`.

```
COLUMNS(ProductId, quantity NUMBER, "Comments")
```

```
COLUMNS(ProductId  VARCHAR2(4000) PATH '$.ProductId',
          quantity   NUMBER          PATH '$.quantity',
          "Comments"  VARCHAR2(4000) PATH '$.Comments')
```

[Example 18-1](#) presents equivalent queries that illustrate this.

If you use **FORMAT JSON** in a given `json_table` column specification then `json_query` semantics are used when projecting the JSON data to the column. The data type specified for the column can be any of the SQL data types that `json_query` can return: `VARCHAR2`, `CLOB`, or `BLOB`.

`json_query` semantics imply the following:

- The projected JSON data is always well-formed. This includes ensuring that non-ASCII characters in string values are escaped as needed. For example, a TAB character (CHARACTER TABULATION, U+0009) is escaped as `\t`.
- `json_query` error handling applies.
- You can use a wrapper clause, to project multiple JSON values as elements in an array.

If you do *not* use `FORMAT JSON` in a given `json_table` column specification then `json_value` semantics are used when projecting the JSON data. The data type specified for the column can be any of the SQL data types that `json_value` can return: `VARCHAR2`, `NUMBER`, `DATE`, `TIMESTAMP`, `TIMESTAMP WITH TIME ZONE`, `SDO_GEOMETRY`, or `CLOB`. `json_value` error handling applies (and you cannot use a wrapper clause).

For example, here the value of column `FirstName` is projected directly using `json_value` semantics, and the value of column `Address` is projected as a JSON string using `json_query` semantics:

```
COLUMNS (FirstName, Address FORMAT JSON)
```

You typically use `FORMAT JSON` when the projected data is a JSON object or array. You typically do not use `FORMAT JSON` when the projected data is a JSON scalar.

Related Topics

- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional **RETURNING** clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no **RETURNING** clause) are described here.
- [Wrapper Clause for SQL/JSON Query Functions JSON_QUERY and JSON_TABLE](#)
SQL/JSON query functions `json_query` and `json_table` accept an optional wrapper clause, which specifies the form of the value returned by `json_query` or used for the data in a `json_table` column. This clause and the default behavior (no wrapper clause) are described here. Examples are provided.
- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.
- [Empty-Field Clause for SQL/JSON Query Functions](#)
SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional **ON EMPTY** clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no **ON EMPTY** clause) are described here.
- [SQL/JSON Function JSON_QUERY](#)
SQL/JSON function `json_query` selects one or more values from JSON data and returns a string (`VARCHAR2`, `CLOB`, or `BLOB` instance) that represents the JSON values. You can thus use `json_query` to retrieve *fragments* of a JSON document.

 **See Also:**

- *Oracle Database SQL Language Reference*
- *Oracle Spatial and Graph Developer's Guide* for information about using Oracle Spatial and Graph data
- GeoJSON.org

18.3 JSON_TABLE Generalizes SQL/JSON Query Functions and Conditions

SQL/JSON function `json_table` generalizes SQL/JSON condition `json_exists` and SQL/JSON functions `json_value` and `json_query`. Everything that you can do using these functions you can do using `json_table`. For the jobs they accomplish, the syntax of these functions is simpler to use than is the syntax of `json_table`.

If you use any of `json_exists`, `json_value`, or `json_query` more than once, or in combination, to access the same data then a single invocation of `json_table` presents the advantage that the data is parsed only once.

Because of this, the optimizer often automatically rewrites multiple invocations of `json_exists`, `json_value` and `json_query` (any combination) to fewer invocations of `json_table` instead, so the data is parsed only once.

[Example 18-3](#) and [Example 18-4](#) illustrate this. They each select the requestor and the set of phones used by each object in column `j_purchaseorder.po_document`. But [Example 18-4](#) parses that column only once, not four times.

Note the following in connection with [Example 18-4](#):

- A JSON value of `null` is a *value* as far as SQL is concerned; it is *not* `NULL`, which in SQL represents the absence of a value (missing, unknown, or inapplicable data). In [Example 18-4](#), if the JSON value of object attribute `zipCode` is `null` then the SQL string `'true'` is returned.
- `json_exists` is a SQL *condition*; you can use it in a SQL `WHERE` clause, a `CASE` statement, or a check constraint. In [Example 18-3](#) it is used in a `WHERE` clause. Function `json_table` employs the semantics of `json_exists` implicitly when you specify keyword `EXISTS`. It must return a SQL *value* in the virtual column. Since Oracle SQL has no Boolean data type, a SQL string `'true'` or `'false'` is used to represent the Boolean value. This is the case in [Example 18-4](#): the `VARCHAR2` value is stored in column `jt.has_zip`, and it is then tested explicitly for equality against the literal SQL string `'true'`.
- JSON field `AllowPartialShipment` has a JSON Boolean value. When `json_value` is applied to that value it is returned as a string. In [Example 18-4](#), data type `VARCHAR2` is used as the column data type. Function `json_table` implicitly uses `json_value` for this column, returning the value as a `VARCHAR2` value, which is then tested for equality against the literal SQL string `'true'`.

Example 18-3 Accessing JSON Data Multiple Times to Extract Data

```

SELECT json_value(po_document, '$.Requestor' RETURNING VARCHAR2(32)),
       json_query(po_document, '$.ShippingInstructions.Phone'
                  RETURNING VARCHAR2(100))
FROM j_purchaseorder
WHERE json_exists(po_document,
                 '$.ShippingInstructions.Address.zipCode')
      AND json_value(po_document, '$.AllowPartialShipment' RETURNING
                     VARCHAR2(5 CHAR))
      = 'true';

```

Example 18-4 Using JSON_TABLE to Extract Data Without Multiple Parses

```

SELECT jt.requestor, jt.phones
FROM j_purchaseorder,
     json_table(po_document, '$'
                COLUMNS (requestor VARCHAR2(32 CHAR) PATH '$.Requestor',
                          phones    VARCHAR2(100 CHAR) FORMAT JSON
                              PATH '$.ShippingInstructions.Phone',
                          partial  VARCHAR2(5 CHAR) PATH
                              '$.AllowPartialShipment',
                          has_zip  VARCHAR2(5 CHAR) EXISTS
                              PATH
                              '$.ShippingInstructions.Address.zipCode')) jt
WHERE jt.partial = 'true' AND jt.has_zip = 'true';

```

Related Topics

- [Using SQL/JSON Function JSON_VALUE With a Boolean JSON Value](#)
JSON has the Boolean values `true` and `false`. When SQL/JSON function `json_value` evaluates a SQL/JSON path expression and the result is JSON `true` or `false`, it can be returned to PL/SQL as a `BOOLEAN` value, or it can be returned to SQL as the `VARCHAR2` value `'true'` or `'false'`.

18.4 Using JSON_TABLE with JSON Arrays

A JSON value can be an array or can include one or more arrays, nested to any number of levels inside other JSON arrays or objects. You can use a `json_table` `NESTED` path clause to project specific elements of an array.

[Example 18-5](#) projects the requestor and associated phone numbers from the JSON data in column `po_document`. The entire JSON array `Phone` is projected as a column of JSON data, `ph_arr`. To format this JSON data as a `VARCHAR2` column, the keywords `FORMAT JSON` are needed.

What if you wanted to project the individual *elements* of JSON array `Phone` and not the array as a whole? [Example 18-6](#) shows one way to do this, which you can use if the array elements are the only data you need to project.

If you want to project both the requestor and the corresponding phone data then the row path expression of [Example 18-6](#) (`$.Phone[*]`) is not appropriate: it targets only the (phone object) elements of array `Phone`.

[Example 18-7](#) shows one way to target both: use a *row path expression* that targets both the name and the entire phones array, and use *column path expressions* that target fields `type` and `number` of individual phone objects.

In [Example 18-7](#) as in [Example 18-5](#), keywords `FORMAT JSON` are needed because the resulting `VARCHAR2` columns contain JSON data, namely arrays of phone types or phone numbers, with one array element for each phone. In addition, unlike the case for [Example 18-5](#), a wrapper clause is needed for column `phone_type` and column `phone_num`, because array `Phone` contains multiple objects with fields `type` and `number`.

Sometimes you might not want the effect of [Example 18-7](#). For example, you might want a column that contains a single phone number (one row per number), rather than one that contains a JSON array of phone numbers (one row for all numbers for a given purchase order).

To obtain that result, you need to tell `json_table` to project the array elements, by using a `json_table NESTED` path clause for the array. A `NESTED` path clause acts, in effect, as an additional row source (row pattern). [Example 18-8](#) illustrates this.

You can use any number of `NESTED` keywords in a given `json_table` invocation.

In [Example 18-8](#) the outer `COLUMNS` clause is the parent of the nested (inner) `COLUMNS` clause. The virtual tables defined are joined using an outer join, with the table defined by the parent clause being the outer table in the join.

(If there were a second `columns` clause nested directly under the same parent, the two nested clauses would be sibling `COLUMNS` clauses.)

Example 18-5 Projecting an Entire JSON Array as JSON Data

```
SELECT jt.*
   FROM j_purchaseorder,
        json_table(po_document, '$'
                  COLUMNS (requestor VARCHAR2(32 CHAR) PATH '$.Requestor',
                             ph_arr   VARCHAR2(100 CHAR) FORMAT JSON
                             PATH '$.ShippingInstructions.Phone')) AS
"JT";
```

Example 18-6 Projecting Elements of a JSON Array

```
SELECT jt.*
   FROM j_purchaseorder,
        json_table(po_document, '$.ShippingInstructions.Phone[*]'
                  COLUMNS (phone_type VARCHAR2(10) PATH '$.type',
                             phone_num  VARCHAR2(20) PATH '$.number')) AS "JT";
```

PHONE_TYPE	PHONE_NUM
Office	909-555-7307
Mobile	415-555-1234

Example 18-7 Projecting Elements of a JSON Array Plus Other Data

```
SELECT jt.*
   FROM j_purchaseorder,
```

```

json_table(po_document, '$'
  COLUMNS (
    requestor VARCHAR2(32 CHAR) PATH '$.Requestor',
    phone_type VARCHAR2(50 CHAR) FORMAT JSON WITH WRAPPER
      PATH '$.ShippingInstructions.Phone[*].type',
    phone_num VARCHAR2(50 CHAR) FORMAT JSON WITH WRAPPER
      PATH '$.ShippingInstructions.Phone[*].number')) AS "JT";

```

REQUESTOR	PHONE_TYPE	PHONE_NUM
-----	-----	-----
Alexis Bull	["Office", "Mobile"]	["909-555-7307", "415-555-1234"]

Example 18-8 JSON_TABLE: Projecting Array Elements Using NESTED

This example shows two equivalent queries that project array elements. The first query uses the simple, dot-notation syntax for the expressions that target the row and column data. The second uses the full syntax.

Except for column `number`, whose SQL identifier is quoted ("`number`"), the SQL column names are, in effect, uppercase. (Column `number` is lowercase.)

In the first query the column names are written exactly the same as the field names that are targeted, including with respect to letter case. Regardless of whether they are quoted, they are interpreted case-sensitively for purposes of establishing the proper path.

The second query has:

- Separate arguments of a JSON column-expression and a SQL/JSON row path-expression
- Explicit column data types of `VARCHAR2(4000)`
- Explicit `PATH` clauses with SQL/JSON column path expressions, to target the object fields that are projected

```

SELECT jt.*
  FROM j_purchaseorder po,
       json_table(po.po_document
                 COLUMNS (Requestor,
                           NESTED ShippingInstructions.Phone[*]
                               COLUMNS (type, "number"))) AS "JT";

```

```

SELECT jt.*
  FROM j_purchaseorder po,
       json_table(po.po_document, '$'
                 COLUMNS (Requestor VARCHAR2(4000) PATH '$.Requestor',
                           NESTED ShippingInstructions.Phone[*]
                               COLUMNS (type VARCHAR2(4000) PATH '$.type',
                                         "number" VARCHAR2(4000) PATH '$.number'))))
 AS "JT";

```

Related Topics

- [Creating a View Over JSON Data Using JSON_TABLE](#)
To improve query performance you can create a view over JSON data that you project to columns using SQL/JSON function `json_table`. To further improve query performance you can create a *materialized view* and place the JSON data *in memory*.
- [SQL/JSON Function JSON_TABLE](#)
SQL/JSON function `json_table` projects specific JSON data to columns of various SQL data types. You use it to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.

18.5 Creating a View Over JSON Data Using JSON_TABLE

To improve query performance you can create a view over JSON data that you project to columns using SQL/JSON function `json_table`. To further improve query performance you can create a *materialized view* and place the JSON data *in memory*.

[Example 18-9](#) defines a view over JSON data. It uses a `NESTED` path clause to project the elements of array `LineItems`.

[Example 18-10](#) defines a materialized view that has the same data and structure as [Example 18-9](#). You cannot update such a materialized view directly; you must treat it as a read-only view and update it indirectly by updating the base table. An error is raised if you try to modify the view directly. If the materialized view is created using keywords `REFRESH` and `ON STATEMENT` then the view is updated automatically whenever you update the base table.

The only differences between [Example 18-9](#) and [Example 18-10](#) are:

- The use of keyword `MATERIALIZED`.
- The use of `BUILD IMMEDIATE`.
- The use of `REFRESH FAST ON STATEMENT WITH PRIMARY KEY`.

The use of `REFRESH FAST` means that the materialized view will be refreshed incrementally. For this to occur, you must use either `WITH PRIMARY KEY` or `WITH ROWID` (if there is no primary key). Oracle recommends that you specify a primary key for a table that has a JSON column and that you use `WITH PRIMARY KEY` when creating a materialized view based on it.

You could use `ON COMMIT` in place of `ON STATEMENT` for the view creation. The former synchronizes the view with the base table only when your table-updating transaction is committed. Until then the table changes are not reflected in the view. If you use `ON STATEMENT` then the view is immediately synchronized after each DML statement. This also means that a view created using `ON STATEMENT` reflects any rollbacks that you might perform. (A subsequent `COMMIT` statement ends the transaction, preventing a rollback.)



See Also:

Oracle Database Data Warehousing Guide

Example 18-9 Creating a View Over JSON Data

```

CREATE OR REPLACE VIEW j_purchaseorder_detail_view
  AS SELECT jt.*
      FROM j_purchaseorder po,
           json_table(po.po_document, '$'
             COLUMNS (
               po_number      NUMBER(10)          PATH '$.PONumber',
               reference       VARCHAR2(30 CHAR)   PATH '$.Reference',
               requestor       VARCHAR2(128 CHAR)  PATH '$.Requestor',
               userid          VARCHAR2(10 CHAR)   PATH '$.User',
               costcenter      VARCHAR2(16)        PATH '$.CostCenter',
               ship_to_name    VARCHAR2(20 CHAR)   PATH '$.ShippingInstructions.name',
               ship_to_street  VARCHAR2(32 CHAR)   PATH
               '$.ShippingInstructions.Address.street',
               ship_to_city    VARCHAR2(32 CHAR)   PATH
               '$.ShippingInstructions.Address.city',
               ship_to_county  VARCHAR2(32 CHAR)   PATH
               '$.ShippingInstructions.Address.county',
               ship_to_postcode VARCHAR2(10 CHAR)  PATH
               '$.ShippingInstructions.Address.postcode',
               ship_to_state   VARCHAR2(2 CHAR)    PATH
               '$.ShippingInstructions.Address.state',
               ship_to_zip     VARCHAR2(8 CHAR)    PATH
               '$.ShippingInstructions.Address.zipCode',
               ship_to_country VARCHAR2(32 CHAR)   PATH
               '$.ShippingInstructions.Address.country',
               ship_to_phone   VARCHAR2(24 CHAR)   PATH
               '$.ShippingInstructions.Phone[0].number',
               NESTED PATH '$.LineItems[*]'
             COLUMNS (
               itemno          NUMBER(38)          PATH '$.ItemNumber',
               description     VARCHAR2(256 CHAR)  PATH
               '$.Part.Description',
               upc_code        VARCHAR2(14 CHAR)   PATH
               '$.Part.UPCCode',
               quantity        NUMBER(12,4)       PATH '$.Quantity',
               unitprice       NUMBER(14,2)       PATH
               '$.Part.UnitPrice')) jt;

```

Example 18-10 Creating a Materialized View Over JSON Data

```

CREATE OR REPLACE MATERIALIZED VIEW j_purchaseorder_materialized_view
  BUILD IMMEDIATE
  REFRESH FAST ON STATEMENT WITH PRIMARY KEY

```

```

AS SELECT jt.*
   FROM j_purchaseorder po,
        json_table(po.po_document, '$'
          COLUMNS (
            po_number          NUMBER(10)          PATH '$.PONumber',
            reference          VARCHAR2(30 CHAR)   PATH '$.Reference',
            requestor         VARCHAR2(128 CHAR)   PATH '$.Requestor',
            userid            VARCHAR2(10 CHAR)    PATH '$.User',
            costcenter        VARCHAR2(16)         PATH '$.CostCenter',
            ship_to_name      VARCHAR2(20 CHAR)
                                PATH '$.ShippingInstructions.name',
            ship_to_street    VARCHAR2(32 CHAR)
                                PATH
'$$.ShippingInstructions.Address.street',
            ship_to_city     VARCHAR2(32 CHAR)
                                PATH
'$$.ShippingInstructions.Address.city',
            ship_to_county   VARCHAR2(32 CHAR)
                                PATH
'$$.ShippingInstructions.Address.county',
            ship_to_postcode VARCHAR2(10 CHAR)
                                PATH
'$$.ShippingInstructions.Address.postcode',
            ship_to_state    VARCHAR2(2 CHAR)
                                PATH
'$$.ShippingInstructions.Address.state',
            ship_to_zip      VARCHAR2(8 CHAR)
                                PATH
'$$.ShippingInstructions.Address.zipCode',
            ship_to_country  VARCHAR2(32 CHAR)
                                PATH
'$$.ShippingInstructions.Address.country',
            ship_to_phone    VARCHAR2(24 CHAR)
                                PATH
'$$.ShippingInstructions.Phone[0].number',
                                NESTED PATH '$.LineItems[*]'
          COLUMNS (
            itemno           NUMBER(38)           PATH '$.ItemNumber',
            description      VARCHAR2(256 CHAR)   PATH
'$$.Part.Description',
            upc_code        VARCHAR2(14 CHAR)    PATH
'$$.Part.UPCCode',
            quantity       NUMBER(12,4)         PATH '$.Quantity',
            unitprice      NUMBER(14,2)         PATH
'$$.Part.UnitPrice')) jt;

```

Related Topics

- [Using JSON_TABLE with JSON Arrays](#)
A JSON value can be an array or can include one or more arrays, nested to any number of levels inside other JSON arrays or objects. You can use a `json_table` NESTED path clause to project specific elements of an array.

Related Topics

- [Using GeoJSON Geographic Data](#)
GeoJSON objects are JSON objects that represent geographic data. Examples are provided of creating GeoJSON data, indexing it, and querying it.

Oracle SQL Function JSON_SERIALIZE

Oracle SQL function `json_serialize` takes JSON data (of any SQL data type, `VARCHAR2`, `CLOB`, or `BLOB`) as input and returns a textual representation of it (as `VARCHAR2`, `CLOB`, or `BLOB` data). `VARCHAR2(4000)` is the default return type.

You typically use `json_serialize` to transform the result of a query. It supports an error clause and a returning clause. You can specify pretty-printing for the result, and you can truncate the result to fit the return type.

Function `json_serialize` always produces JSON data that conforms to the JSON standard (RFC 4627). The returned data uses only the *standard* data types of the JSON language: object, array, and the scalar types string, number, Boolean, and null.

You can use `json_serialize` to convert binary JSON data to textual form (`CLOB` or `VARCHAR2`), or to transform textual JSON data by pretty-printing it or escaping non-ASCII Unicode characters in it. An important use case is serializing JSON data that is stored in a `BLOB` column. A `BLOB` result is in the AL32UTF8 character set. But whatever the data type returned by `json_serialize`, the returned data represents textual JSON data.

Note:

The default error handling for `json_serialize` is `NULL ON ERROR`, which means that when an error occurs a SQL `NULL` value is returned — an error is not raised.

In particular, be aware of this in the context of returning serialized data that is too large. If you are unsure of the return size, use `RETURNING CLOB` or `RETURNING BLOB`, to allow for more than the default of 4000 bytes of `VARCHAR2` data. Use clause `ERROR ON ERROR` if you want to be sure to raise an error when the result is too large.

See Also:

Oracle Database SQL Language Reference for information about Oracle SQL function `json_serialize`

Example 19-1 Using JSON_SERIALIZE To Convert BLOB Data To Pretty-Printed Text

This example serializes and pretty-prints the JSON purchase order that has 1600 as the value of field `PONumber` data, which is selected from column `po_document` of table `j_purchaseorder`. The return-value data type is `VARCHAR2(4000)` (the default return type).

[Example 4-2](#) shows the insertion of such purchase-order data into a `VARCHAR2` table. But `json_serialize` is especially useful for serializing `BLOB` data. See [Example 9-1](#) for how to create a table with a `BLOB` column of JSON data.

```
SELECT json_serialize(po_document PRETTY) FROM j_purchaseorder;
```

Related Topics

- [RETURNING Clause for SQL Query Functions](#)
SQL functions `json_value`, `json_query`, `json_serialize`, and `json_mergepatch` accept an optional `RETURNING` clause, which specifies the data type of the value returned by the function. This clause and the default behavior (no `RETURNING` clause) are described here.
- [Error Clause for SQL Query Functions and Conditions](#)
Some SQL query functions and conditions accept an optional error clause, which specifies handling for a runtime error that is raised by the function or condition. This clause and the default behavior (no error clause) are summarized here.

JSON Data Guide

A JSON data guide lets you discover information about the structure and content of JSON documents stored in Oracle Database.

Some ways that you can use this information include:

- Generating a JSON Schema document that describes the set of JSON documents.
- Creating views that you can use to perform SQL operations on the data in the documents.
- Automatically adding or updating virtual columns that correspond to added or changed fields in the documents.
- [Overview of JSON Data Guide](#)
A data guide is a summary of the structural and type information contained in a set of JSON documents. It records metadata about the fields used in those documents.
- [Persistent Data-Guide Information: Part of a JSON Search Index](#)
JSON data-guide information can be saved persistently as part of the JSON search index infrastructure, and this information is updated automatically as new JSON content is added. This is the case by default, when you create a JSON search index: data-guide information is part of the index infrastructure.
- [Data-Guide Formats and Ways of Creating a Data Guide](#)
There are two formats for a data guide: flat and hierarchical. Both are made available to SQL and PL/SQL as CLOB data. You can construct a data guide from the data-guide information stored in a JSON search index or by scanning JSON documents.
- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.
- [Data-Dictionary Views For Persistent Data-Guide Information](#)
You can query static data-dictionary views to see which tables have JSON columns with data guide-enabled JSON search indexes and to extract JSON object field information that is recorded in dataguide-enabled JSON search indexes.
- [Specifying a Preferred Name for a Field Column](#)
You can project JSON fields from your data as columns in a database view or as virtual columns added to the same table that contains the JSON column. You can specify a preferred name for such a column.
- [Creating a View Over JSON Data Based on Data-Guide Information](#)
Based on data-guide information, you can create a database view whose columns project particular scalar fields present in a set of JSON documents. You can choose the fields to project by editing a hierarchical data guide or by specifying a SQL/JSON path expression and a minimum frequency of field occurrence.

- [Adding and Dropping Virtual Columns For JSON Fields Based on Data-Guide Information](#)
Based on data-guide information for a JSON column, you can project scalar fields from that JSON data as virtual columns in the same table. The scalar fields projected are those that are not under an array.
- [Change Triggers For Data Guide-Enabled Search Index](#)
When JSON data changes, some information in a data guide-enabled JSON search index is automatically updated. You can specify a procedure whose invocation is triggered whenever this happens. You can define your own PL/SQL procedure for this, or you can use the predefined change-trigger procedure `add_vc`.
- [Multiple Data Guides Per Document Set](#)
A data guide reflects the shape of a given set of JSON documents. If a JSON column contains different types of documents, with different structure or type information, you can create and use different data guides for the different kinds of documents.
- [Querying a Data Guide](#)
A data guide is information about a set of JSON documents. You can query it from a flat data guide that you obtain using either Oracle SQL function `json_dataguide` or PL/SQL function `DBMS_JSON.get_index_dataguide`. In the latter case, a data guide-enabled JSON search index must be defined on the JSON data.
- [A Flat Data Guide For Purchase-Order Documents](#)
The fields of a sample flat data guide are described. It corresponds to a set of purchase-order documents.
- [A Hierarchical Data Guide For Purchase-Order Documents](#)
The fields of a sample hierarchical data guide are described. It corresponds to a set of purchase-order documents.

 **See Also:**

- *JSON Schema: core definitions and terminology [json-schema-core](#)*
- *JSON Schema: interactive and non interactive validation*

20.1 Overview of JSON Data Guide

A data guide is a summary of the structural and type information contained in a set of JSON documents. It records metadata about the fields used in those documents.

For example, for the JSON object presented in [Example 2-1](#), a data guide specifies that the document has, among other things, an object `ShippingInstructions` with fields `name`, `Address`, and `Phone`, of types `string`, `object`, and `array`, respectively. The structure of object `Address` is recorded similarly, as are the types of the elements in array `Phone`.

JSON data-guide information can be saved persistently as part of the JSON search index infrastructure, and this information is updated automatically as new JSON content is added. This is the case by default, when you create a JSON search index: data-guide information is part of the index infrastructure.

You can use a data guide:

- As a basis for developing applications that involve data mining, business intelligence, or other analysis of JSON documents.
- As a basis for providing user assistance about requested JSON information, including search.
- To check or manipulate new JSON documents before adding them to a document set (for example: validate, type-check, or exclude certain fields).

For such purposes you can:

- Query a data guide directly for information about the document set, such as field lengths or which fields occur with at least a certain frequency.
- Create views, or add virtual columns, that project particular JSON fields of interest, based on their significance according to a data guide.

 **Note:**

- The advantages of virtual columns over a view are that you can build an index on a virtual column and you can obtain statistics on it for the optimizer.
- Virtual columns, like columns in general, are subject to the 1000-column limit for a given table.

Related Topics

- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.
- [JSON Search Index: Ad Hoc Queries and Full-Text Search](#)
A JSON search index is a *general* index. It can improve the performance of both (1) ad hoc structural queries, that is, queries that you might not anticipate or use regularly, and (2) full-text search. It is an Oracle Text index that is designed specifically for use with JSON data.
- [Querying a Data Guide](#)
A data guide is information about a set of JSON documents. You can query it from a flat data guide that you obtain using either Oracle SQL function `json_dataguide` or PL/SQL function `DBMS_JSON.get_index_dataguide`. In the latter case, a data guide-enabled JSON search index must be defined on the JSON data.
- [Creating a View Over JSON Data Based on a Hierarchical Data Guide](#)
You can use a hierarchical data guide to create a database view whose columns project specified JSON fields from your documents. The fields projected are those in the data guide. You can edit the data guide to include only the fields that you want to project.
- [Adding and Dropping Virtual Columns For JSON Fields Based on Data-Guide Information](#)
Based on data-guide information for a JSON column, you can project scalar fields from that JSON data as virtual columns in the same table. The scalar fields projected are those that are not under an array.

 **See Also:**

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database SQL Language Reference* for information about SQL function `json_dataguide`

20.2 Persistent Data-Guide Information: Part of a JSON Search Index

JSON data-guide information can be saved persistently as part of the JSON search index infrastructure, and this information is updated automatically as new JSON content is added. This is the case by default, when you create a JSON search index: data-guide information is part of the index infrastructure.

You can use `CREATE SEARCH INDEX` with keywords `FOR JSON` to create a search index, a data guide, or both at the same time. The default behavior is to create both.

To create persistent data-guide information as part of a JSON search index *without enabling support for search* in the index, you specify `SEARCH_ON NONE` in the `PARAMETERS` clause for `CREATE SEARCH INDEX`, but you leave the value for `DATAGUIDE` as `ON` (the default value). [Example 20-1](#) illustrates this.

You can use `ALTER INDEX ... REBUILD` to enable or disable data-guide support for an *existing* JSON search index. [Example 20-2](#) illustrates this — it disables the data-guide support that is added by default in [Example 26-17](#).

 **Note:**

To create a data guide-enabled JSON search index, or to data guide-enable an existing JSON search index, you need database privilege `CTXAPP` and Oracle Database Release 12c (12.2.0.1) or later.

 **Note:**

A data guide-enabled JSON search index can be built only on a column that has an `is_json` check constraint. Furthermore, for the data-guide information in the index to be updated the check constraint must be enabled.

If the check constraint becomes disabled for some reason then you must *rebuild the data-guide information* in the index and *re-enable the check constraint*, to resume automatic data-guide support updating, as follows:

```
ALTER INDEX index_name REBUILD ('dataguide off');  
ALTER INDEX index_name REBUILD ('dataguide on');  
ALTER TABLE table_name ENABLE CONSTRAINT  
is_json_check_constraint_name;
```

In particular, using SQL*Loader (`sqlldr`) disables `is_json` check constraints.

Because persistent data-guide information is part of the search index infrastructure, it is always *live*: its content is automatically updated whenever the index is synchronized. Changes in the indexed data are reflected in the search index, including in its data-guide information, only after the index is synchronized.

In addition, update of data-guide information in a search index is always *additive*: none of it is ever deleted. This is another reason that the index often does not accurately reflect the data in its document set: deletions within the documents it applies to are *not* reflected in its data-guide information. If you need to ensure that such information accurately reflects the current data then you must drop the JSON search index and create it anew.

The persistent data-guide information in a search index can also include *statistics*, such as how frequently each JSON field is used in the document set. Statistics are present only if you explicitly gather them on the document set (gather them on the JSON search index, for example). They are not updated automatically — gather statistics anew if you want to be sure they are up to date. [Example 20-3](#) gathers statistics on the JSON data indexed by JSON search index `po_search_idx`, which is created in [Example 26-17](#).

 **Note:**

When a local data guide-enabled JSON search index is created in a *sharding* environment, each individual shard contains the data-guide information for the JSON documents stored in that shard. For this reason, if you invoke data guide-related operations on the shard *catalog* database then they will have no effect.

Considerations for a Data Guide-Enabled Search Index on a Partitioned Table

The data-guide information in a data guide-enabled JSON search index that is local to a partitioned table is not partitioned. It is shared among all partitions.

Because the data-guide information in the index is only additive, dropping, merging, splitting, or truncating partitions has no impact on the index.

Exchanging a partitioned table with a table that is not partitioned updates the data-guide information in an index on the partitioned table, but any data guide-enabled index on the non-partitioned table must be rebuilt.

Avoid Persistent Data-Guide Information If Serializing Hash-Table Data

If you serialize Java hash tables or associative arrays (such as are found in JavaScript) as JSON objects, then avoid the use of persistent data-guide information.

The default hash-table serialization provided by popular libraries such as GSON and Jackson produces textual JSON documents with object field names that are taken from the hash-table key entries and with field values taken from the corresponding Java hash-table values. Serializing a single Java hash-table entry produces a new (unique) JSON field and value.

Persisted data-guide information reflects the shape of your data, and it is updated automatically as new JSON documents are inserted. Each hash-table key–value pair results in a separate entry in the JSON search index. Such serialization can thus greatly increase the size of the information maintained in the index. In addition to the large size, the many index updates affect performance negatively, making DML slow.

If you serialize a hash table or an associative array instead as a JSON array of objects, each of which includes a field derived from a hash-table key entry, then there are no such problems.

The default serialization of a hash table or associative array as a JSON object is indistinguishable from an object that has field names assigned by a developer. A JSON data guide cannot tell which (metadata-like) field names have been assigned by a developer and which (data-like) names might have been derived from a hash table or associative array. It treats all field names as essentially metadata, as if specified by a developer.

For example:

- If you construct an application object using a hash table that has `animalName` as the hash key and sets of animal properties as values then the resulting default serialization is a single JSON object that has a *separate field* ("`cat`", "`mouse`",...) for each hash-table entry, with the field value being an object with the corresponding animal properties. This can be problematic in terms of data-guide size and performance because of the typically large number of fields ("`cat`", "`mouse`",...) derived from the hash key.
- If you instead construct an application array of `animal` structures, each of which has a field `animalName` (with value "`cat`" or "`mouse`"...) then the resulting serialization is a JSON array of objects, each of which has the same field, `animalName`. The corresponding data guide has no size or performance problem.

Example 20-1 Enabling Persistent Support for a JSON Data Guide But Not For Search

```
CREATE SEARCH INDEX po_dg_only_idx ON j_purchaseorder (po_document) FOR
JSON
PARAMETERS ('SEARCH_ON NONE');
```

Example 20-2 Disabling JSON Data-Guide Support For an Existing JSON Search Index

```
ALTER INDEX po_search_idx REBUILD PARAMETERS ('DATAGUIDE OFF');
```

Example 20-3 Gathering Statistics on JSON Data Using a JSON Search Index

```
EXEC DBMS_STATS.gather_index_stats(docuser, po_search_idx, NULL, 99);
```

Related Topics

- [JSON Search Index: Ad Hoc Queries and Full-Text Search](#)
A JSON search index is a *general* index. It can improve the performance of both (1) ad hoc structural queries, that is, queries that you might not anticipate or use regularly, and (2) full-text search. It is an Oracle Text index that is designed specifically for use with JSON data.

 **See Also:**

- *Oracle Text Reference* for information about the `PARAMETERS` clause for `CREATE SEARCH INDEX`
- *Oracle Text Reference* for information about the `PARAMETERS` clause for `ALTER INDEX ... REBUILD`
- *Faster XML / Jackson* for information about the Jackson JSON processor
- *google / gson* for information about the GSON Java library

20.3 Data-Guide Formats and Ways of Creating a Data Guide

There are two formats for a data guide: flat and hierarchical. Both are made available to SQL and PL/SQL as `CLOB` data. You can construct a data guide from the data-guide information stored in a JSON search index or by scanning JSON documents.

- You can use a *flat* data guide to *query* data-guide information such as field frequencies and types.

A flat data guide is represented in JSON as an *array* of objects, each of which represents the JSON data of a specific *path* in the document set. [A Flat Data Guide For Purchase-Order Documents](#) describes a flat data guide for the purchase-order data of [Example 2-1](#).

- You can use a *hierarchical* data guide to create a view, or to add virtual columns, using particular fields that you choose on the basis of data-guide information.

A hierarchical data guide is represented in JSON as an *object* with nested JSON data, in the same format as that defined by JSON Schema (version 4, `json-schema-core`). [A Hierarchical Data Guide For Purchase-Order Documents](#) describes a hierarchical data guide for the purchase-order data of [Example 2-1](#).

You use PL/SQL function `DBMS_JSON.get_index_dataguide` to obtain a data guide from the data-guide information stored in a JSON search index.

You can also use SQL aggregate function `json_dataguide` to scan your document set and construct a data guide for it, whether or not it has a data guide-enabled search index. The data guide accurately reflects the document set at the moment of function invocation.

Table 20-1 SQL and PL/SQL Functions to Obtain a Data Guide

Uses Data Guide-Enabled Search Index?	Flat Data Guide	Hierarchical Data Guide
Yes	PL/SQL function <code>get_index_dataguide</code> with format <code>DBMS_JSON.FORMAT_FLAT</code>	PL/SQL function <code>get_index_dataguide</code> with format <code>DBMS_JSON.FORMAT_HIERARCHICAL</code>
No	SQL function <code>json_dataguide</code> , with no format argument or with <code>DBMS_JSON.FORMAT_FLAT</code> as the format argument	SQL function <code>json_dataguide</code> , with <code>DBMS_JSON.FORMAT_HIERARCHICAL</code> as the format argument

Advantages of obtaining a data guide based on a data guide-enabled JSON search index include:

- Additive updates to the document set are automatically reflected in the persisted data-guide information whenever the index is synced.
- Because this data-guide information is persisted, obtaining a data guide based on it (using PL/SQL function `get_index_dataguide`) is faster than obtaining a data guide by analyzing the document set (using SQL function `json_dataguide`).
- If you have gathered statistics on the document set then these are included in the stored information and in a data guide obtained from it.
- Column-name conflicts encountered when creating a view or virtual columns are automatically resolved.

Advantages of obtaining a data guide without using a data guide-enabled JSON search index include assurance that the data guide is accurate and the lack of index maintenance overhead. In addition, a data guide that is not derived from an index is appropriate in some particular use cases:

- The JSON data is in an external table. You cannot create an index on it.
- The JSON column could be indexed, but the index would not be very useful. This can be the case, for example, if the column contains different kinds of documents. In this case, it can sometimes be helpful to add a column to the table that identifies the kind of document stored in the JSON column. You can then use the data guide with SQL aggregate functions and `GROUP BY`. See [Multiple Data Guides Per Document Set](#).

Related Topics

- [A Flat Data Guide For Purchase-Order Documents](#)
The fields of a sample flat data guide are described. It corresponds to a set of purchase-order documents.

- [A Hierarchical Data Guide For Purchase-Order Documents](#)
The fields of a sample hierarchical data guide are described. It corresponds to a set of purchase-order documents.
- [Persistent Data-Guide Information: Part of a JSON Search Index](#)
JSON data-guide information can be saved persistently as part of the JSON search index infrastructure, and this information is updated automatically as new JSON content is added. This is the case by default, when you create a JSON search index: data-guide information is part of the index infrastructure.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database SQL Language Reference* for information about SQL function `json_dataguide`
- *Oracle Database SQL Language Reference* for information about PL/SQL constants `DBMS_JSON.FORMAT_FLAT` and `DBMS_JSON.FORMAT_HIERARCHICAL`

20.4 JSON Data-Guide Fields

The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.

A given occurrence of a field in a data guide corresponds to a field that is present in one or more JSON documents of the document set.

JSON Schema Fields (Keywords)

A JSON Schema is a JSON document that contains a JSON object, which can itself contain child objects (subschemas). Fields that are defined by JSON Schema are called JSON Schema **keywords**. [Table 20-2](#) describes the keywords that can be used in an Oracle JSON data guide. Keywords `properties`, `items`, and `oneOf` are used only in a hierarchical JSON data guide (which is a JSON schema). Keyword `type` is used in both flat and hierarchical data guides.

Table 20-2 JSON Schema Fields (Keywords)

Field (Keyword)	Value Description
<code>properties</code>	An object whose members represent the properties of a JSON object used in JSON data that is represented by the hierarchical data guide (JSON schema).
<code>items</code>	An object whose members represent the elements (items) of an array used in JSON data represented by the hierarchical data guide (JSON schema).
<code>oneOf</code>	An array, each of whose items represents one or more occurrences of a JSON field in the JSON data represented by the hierarchical data guide (JSON schema).

Table 20-2 (Cont.) JSON Schema Fields (Keywords)

Field (Keyword)	Value Description
type	A string naming the type of some JSON data represented by the (flat or hierarchical) data guide. The possible values are: "number", "string", "boolean", "null", "object", "array", and "GeoJSON".

Oracle-Specific JSON Data-Guide Fields

In addition to JSON Schema keywords, a JSON data guide can contain Oracle data guide-specific fields. The field names all have the prefix `o:`. They are described in [Table 20-3](#).

Table 20-3 Oracle-Specific Data-Guide Fields

Field	Value Description
<code>o:path</code>	Path through the JSON documents to the JSON field. Used only in a <i>flat</i> data guide. The value is a simple SQL/JSON path expression (no filter expression), possibly with relaxation (implicit array wrapping and unwrapping), and possibly with a wildcard array step. It has no array steps with array indexes or range specifications, and it has no function step. See SQL/JSON Path Expression Syntax .
<code>o:length</code>	Maximum length of the JSON field value, in bytes. The value is always a power of two. For example, if the maximum length of all actual field values is 5 then the value of <code>o:length</code> is 8, the smallest power of two greater than or equal to 5.
<code>o:preferred_column_name</code>	An identifier, case-sensitive and unique to a given data guide, that you prefer as the name to use for a view column or a virtual column that is created using the data guide. This field is absent if the data guide was obtained using SQL function <code>json_dataguide</code> with format parameter <code>DBMS_JSON.FORMAT_FLAT</code> or without any format parameter (<code>DBMS_JSON.FORMAT_FLAT</code> is the default).
<code>o:frequency</code>	Percentage of JSON documents that contain the given field. Duplicate occurrences of a field under the same array are ignored. (Available only if statistics were gathered on the document set.) This field is absent if the data guide was obtained using SQL function <code>json_dataguide</code> .
<code>o:num_nulls</code>	Number of documents whose value for the targeted scalar field is JSON <code>null</code> . (Available only if statistics were gathered on the document set.) This field is absent if the data guide was obtained using SQL function <code>json_dataguide</code> .
<code>o:high_value</code>	Highest value for the targeted scalar field, among all documents. (Available only if statistics were gathered on the document set.) This field is absent if the data guide was obtained using SQL function <code>json_dataguide</code> .

Table 20-3 (Cont.) Oracle-Specific Data-Guide Fields

Field	Value Description
<code>o:low_value</code>	Lowest value for the targeted scalar field, among all documents. (Available only if statistics were gathered on the document set.) This field is absent if the data guide was obtained using SQL function <code>json_dataguide</code> .
<code>o:last_analyzed</code>	Date and time when statistics were last gathered on the document set. (Available only if statistics were gathered on the document set.) This field is absent if the data guide was obtained using SQL function <code>json_dataguide</code> .

The data-guide information for documents that contain a JSON array with *only scalar* elements records the path and type for both (1) the array and (2) all of the array elements taken together. For the *elements*:

- The `o:path` value is the `o:path` value for the array, followed by an array with a wildcard (`[*]`), which indicates all array elements.
- The `type` value is the `type` string, if the scalar types are *not the same* for all elements in all documents. If *all* of the scalar elements the array have the *same* type, across *all* documents, then that type is recorded.

For example, if, in *all* documents, *all* of the elements in the array value for object field `serial_numbers` are JSON numbers, then `type` for the array elements is `number`. Otherwise it is `string`.

When present, the default value of field `o:preferred_column_name` depends on whether the data guide was obtained using SQL function `json_dataguide` (with format `DBMS_JSON.FORMAT_HIERARCHICAL`) or using PL/SQL function `DBMS_JSON.get_index_dataguide`:

- `json_dataguide` (hierarchical format) — Same as the corresponding JSON field name.
- `get_index_dataguide` — Same as the corresponding JSON field name, prefixed with the JSON column name followed by `$`, and with any non-ASCII characters removed. If the resulting field name already exists in the same data guide then it is suffixed with a new sequence number, to make it unique.

The column name is uppercase unless the column was defined using escaped lowercase letters (for example, `'PO_Column'` instead of `po_column`).

For example, the default value for field `User` for data in JSON column `po_document` is `PO_DOCUMENT$User`.

You can use PL/SQL procedure `DBMS_JSON.rename_column` to set the value of `o:preferred_column_name` for a given field and type.

Field `o:preferred_column_name` is used to name a new, virtual column in the table that contains the JSON column, or it is used to name a column in a new view that also contains the other columns of the table. In either case, the name specified by `o:preferred_column_name` must be *unique* with respect to the other column names of the table. In addition, the name must be *unique* across all JSON fields of any type in

the document set. When you use `DBMS_JSON.get_index_dataguide`, the default name is *guaranteed* to be unique in these ways.

If the name you specify with `DBMS_JSON.rename_column` causes a name conflict then the specified name is ignored and a system-generated name is used instead.

Related Topics

- [Specifying a Preferred Name for a Field Column](#)
You can project JSON fields from your data as columns in a database view or as virtual columns added to the same table that contains the JSON column. You can specify a preferred name for such a column.
- [A Flat Data Guide For Purchase-Order Documents](#)
The fields of a sample flat data guide are described. It corresponds to a set of purchase-order documents.
- [A Hierarchical Data Guide For Purchase-Order Documents](#)
The fields of a sample hierarchical data guide are described. It corresponds to a set of purchase-order documents.
- [Using GeoJSON Geographic Data](#)
GeoJSON objects are JSON objects that represent geographic data. Examples are provided of creating GeoJSON data, indexing it, and querying it.



See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`
- *Oracle Database SQL Language Reference* for information about SQL function `json_dataguide`
- *Oracle Spatial and Graph Developer's Guide* for information about using GeoJSON data with Oracle Spatial and Graph
- *Oracle Spatial and Graph Developer's Guide* for information about Oracle Spatial and Graph and `SDO_GEOMETRY` object type
- GeoJSON.org for information about GeoJSON

20.5 Data-Dictionary Views For Persistent Data-Guide Information

You can query static data-dictionary views to see which tables have JSON columns with data guide-enabled JSON search indexes and to extract JSON object field information that is recorded in dataguide-enabled JSON search indexes.

Tables that do not have JSON columns with data guide-enabled indexes are not present in the views.

You can use the following views to *find columns* that have data guide-enabled JSON search indexes. The views have columns `TABLE_NAME` (the table name), `COLUMN_NAME` (the JSON column name), and `DATAGUIDE` (a data guide).

- `USER_JSON_DATAGUIDES` — tables owned by the current user
- `ALL_JSON_DATAGUIDES` — tables accessible by the current user
- `DBA_JSON_DATAGUIDES` — all tables

If the JSON column has a data guide-enabled JSON search index then the value of column `DATAGUIDE` is the data guide for the JSON column, in flat format as a `CLOB` instance. If it does not have a data guide-enabled index then there is no row for that column in the view.

You can use the following views to extract JSON field path and type information that is recorded in dataguide-enabled JSON search indexes. The views have columns `TABLE_NAME`, `COLUMN_NAME`, `PATH`, `TYPE`, and `LENGTH`. Columns `PATH`, `TYPE`, and `LENGTH` correspond to the values of data-guide fields `o:path`, `o:type`, and `o:length`, respectively.

- `USER_JSON_DATAGUIDE_FIELDS` — tables owned by the current user
- `ALL_JSON_DATAGUIDE_FIELDS` — tables accessible by the current user
- `DBA_JSON_DATAGUIDE_FIELDS` — all tables

In the case of both types of view, a view whose name has the prefix `ALL_` or `DBA_` also has column `OWNER`, whose value is the table owner.

See Also:

- *Oracle Database Reference* for information about `ALL_JSON_DATAGUIDES` and the related data-dictionary views
- *Oracle Database Reference* for information about `ALL_JSON_DATAGUIDE_FIELDS` and the related data-dictionary views

20.6 Specifying a Preferred Name for a Field Column

You can project JSON fields from your data as columns in a database view or as virtual columns added to the same table that contains the JSON column. You can specify a preferred name for such a column.

A data guide obtained from the same document set is used to define this projection. The name of each projected column is taken from data-guide field `o:preferred_column_name` for the JSON data field to be projected. Specifying your preferred name changes the value of this data-guide field.

If your JSON data has a data guide-enabled search index then you can use procedure `DBMS_JSON.rename_column` to specify your preferred name for the column projected from a given field. [Example 20-4](#) illustrates this. It specifies preferred names for the columns to be projected from various fields, as described in [Table 20-4](#). (The fields are projected as columns when you use procedure `DBMS_JSON.create_view`, `DBMS_JSON.create_view_on_path`, or `DBMS_JSON.add_virtual_columns`.)

Table 20-4 Preferred Names for Some JSON Field Columns

Field	JSON Type	Preferred Column Name
PONumber	number	PONumber
Phone (phone as string, not object – just the number)	string	Phone
type (phone type)	string	PhoneType
number (phone number)	string	PhoneNumber
ItemNumber (line-item number)	number	ItemNumber
Description (part description)	string	PartDescription

 **See Also:**

- [JSON Data-Guide Fields](#) for information about the default value of field `o:preferred_column_name` and the possibility of name conflicts when you use `DBMS_JSON.rename_column`
- [Creating a Table With a JSON Column](#) for information about the JSON data referenced here
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view_on_path`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.add_virtual_columns`

Example 20-4 Specifying Preferred Column Names For Some JSON Fields

```

BEGIN
  DBMS_JSON.rename_column(
    'J_PURCHASEORDER', 'PO_DOCUMENT', '$.PONumber',
    DBMS_JSON.TYPE_NUMBER, 'PONumber');
  DBMS_JSON.rename_column(
    'J_PURCHASEORDER', 'PO_DOCUMENT', '$.ShippingInstructions.Phone',
    DBMS_JSON.TYPE_STRING, 'Phone');
  DBMS_JSON.rename_column(
    'J_PURCHASEORDER', 'PO_DOCUMENT',
    '$.ShippingInstructions.Phone.type',
    DBMS_JSON.TYPE_STRING, 'PhoneType');
  DBMS_JSON.rename_column(
    'J_PURCHASEORDER', 'PO_DOCUMENT',
    '$.ShippingInstructions.Phone.number',
    DBMS_JSON.TYPE_STRING, 'PhoneNumber');
  DBMS_JSON.rename_column(
    'J_PURCHASEORDER', 'PO_DOCUMENT', '$.LineItems.ItemNumber',

```

```

DBMS_JSON.TYPE_NUMBER, 'ItemNumber');
DBMS_JSON.rename_column(
  'J_PURCHASEORDER', 'PO_DOCUMENT', '$.LineItems.Part.Description',
  DBMS_JSON.TYPE_STRING, 'PartDescription');
END;
/

```

20.7 Creating a View Over JSON Data Based on Data-Guide Information

Based on data-guide information, you can create a database view whose columns project particular scalar fields present in a set of JSON documents. You can choose the fields to project by editing a hierarchical data guide or by specifying a SQL/JSON path expression and a minimum frequency of field occurrence.

(You can create multiple views based on the same JSON document set, projecting different fields. See [Multiple Data Guides Per Document Set](#).)

You can create a view by projecting JSON fields using SQL/JSON function `json_table` — see [Creating a View Over JSON Data Using JSON_TABLE](#). An alternative is to use PL/SQL procedure `DBMS_JSON.create_view` or `DBMS_JSON.create_view_on_path` to create a view by projecting fields that you choose based on available data-guide information.

This information can come from either a hierarchical data guide that includes only the fields to project or from a data guide-enabled JSON search index together with a SQL/JSON path expression and a minimum field frequency.

In the former case, use procedure `create_view`. You can edit a (hierarchical) data guide to specify the fields you want. In this case you do *not* need a data guide-enabled search index.

In the latter case, use procedure `create_view_on_path`. In this case you need a data guide-enabled search index, but you do not need a data guide. You provide a SQL/JSON path expression and possibly a minimum frequency of occurrence. The fields in the document set that are projected include both:

- All scalar fields that are not under an array.
- All scalar fields present, at any level, in the data that is targeted by a given SQL/JSON path expression.

Regardless of which way you create the view, in addition to the JSON fields that are projected as columns, the non-JSON columns of the table are also columns of the view.

If you use procedure `create_view_on_path` then the `PATH` argument you provide must be a simple SQL/JSON path expression (no filter expression), possibly with relaxation (implicit array wrapping and unwrapping), but with no array steps and no function step. See [SQL/JSON Path Expression Syntax](#).

However it is created, the data guide that serves as the basis for a given view definition is static and does not necessarily faithfully reflect the current data in the document set. The fields that are projected for the view are determined when the view is created.

In particular, if you use `create_view_on_path` (which requires a data guide-enabled search index) then what counts are the fields specified by the given path expression and that have at least the given frequency, based on the *index data at the time of the view creation*.

- [Creating a View Over JSON Data Based on a Hierarchical Data Guide](#)
You can use a hierarchical data guide to create a database view whose columns project specified JSON fields from your documents. The fields projected are those in the data guide. You can edit the data guide to include only the fields that you want to project.
- [Creating a View Over JSON Data Based on a Path Expression](#)
You can use the information in a data guide-enabled JSON search index to create a database view whose columns project JSON fields from your documents. The fields projected are the scalar fields not under an array plus the scalar fields in the data targeted by a specified SQL/JSON path expression.

Related Topics

- [Creating a View Over JSON Data Using JSON_TABLE](#)
To improve query performance you can create a view over JSON data that you project to columns using SQL/JSON function `json_table`. To further improve query performance you can create a *materialized view* and place the JSON data *in memory*.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view_on_path`

20.7.1 Creating a View Over JSON Data Based on a Hierarchical Data Guide

You can use a hierarchical data guide to create a database view whose columns project specified JSON fields from your documents. The fields projected are those in the data guide. You can edit the data guide to include only the fields that you want to project.

You can obtain a hierarchical data guide using either PL/SQL function `DBMS_JSON.get_index_dataguide` or SQL function `json_dataguide` with argument `DBMS_JSON.FORMAT_HIERARCHICAL`. In the former case a data guide-enabled JSON search index must be defined on the column of JSON data.

You can edit the data guide obtained to include only specific fields, change the length of given types, or rename fields. The resulting data guide specifies which fields of the JSON data to project as columns of the view.

You use PL/SQL procedure `DBMS_JSON.create_view` to create the view.

Example 20-5 illustrates this using a data guide obtained with `DBMS_JSON.get_index_dataguide`. **Example 20-6** illustrates it using a data guide obtained with `json_dataguide` with argument `DBMS_JSON.FORMAT_HIERARCHICAL`.

If you create a view using the data guide obtained with `json_dataguide` then GeoJSON data in your documents is supported. In this case the view column corresponding to the GeoJSON data has SQL data type `SDO_GEOMETRY`. For that you pass constant `DBMS_JSON.GEOJSON` or `DBMS_JSON.GEOJSON+DBMS_JSON.PRETTY` as the third argument to `json_dataguide`.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`
- *Oracle Database SQL Language Reference* for information about SQL function `json_dataguide`
- *Oracle Database SQL Language Reference* for information about PL/SQL constant `DBMS_JSON.FORMAT_HIERARCHICAL`

Example 20-5 Creating a View Using a Hierarchical Data Guide Obtained With `GET_INDEX_DATAGUIDE`

This example creates a view that projects all of the fields present in the hierarchical data guide that is obtained from the data guide-enabled JSON search index on JSON column `po_document` of table `j_purchaseorder`. (Columns whose names are *italic* in the `describe` command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`.)

```
EXEC DBMS_JSON.create_view(
    'VIEW1',
    'J_PURCHASEORDER',
    'PO_DOCUMENT',
    DBMS_JSON.get_index_dataguide('J_PURCHASEORDER',
                                  'PO_DOCUMENT',
                                  DBMS_JSON.FORMAT_HIERARCHICAL));
```

```
DESCRIBE view1
```

Name	Null?	Type
DATE_LOADED		TIMESTAMP(6) WITH TIME ZONE
ID	NOT NULL	RAW(16)
<i>PO_DOCUMENT\$User</i>		VARCHAR2(8)
<i>PONumber</i>		NUMBER
<i>PO_DOCUMENT\$Reference</i>		VARCHAR2(16)
<i>PO_DOCUMENT\$Requestor</i>		VARCHAR2(16)
<i>PO_DOCUMENT\$CostCenter</i>		VARCHAR2(4)

PO_DOCUMENT\$AllowPartialShipment	VARCHAR2(4)
PO_DOCUMENT\$name	VARCHAR2(16)
Phone	VARCHAR2(16)
PO_DOCUMENT\$city	VARCHAR2(32)
PO_DOCUMENT\$state	VARCHAR2(2)
PO_DOCUMENT\$street	VARCHAR2(32)
PO_DOCUMENT\$country	VARCHAR2(32)
PO_DOCUMENT\$zipCode	NUMBER
PO_DOCUMENT\$SpecialInstructions	VARCHAR2(8)
PO_DOCUMENT\$UPCCCode	NUMBER
PO_DOCUMENT\$UnitPrice	NUMBER
PartDescription	VARCHAR2(32)
PO_DOCUMENT\$Quantity	NUMBER
ItemNumber	NUMBER
PhoneType	VARCHAR2(8)
PhoneNumber	VARCHAR2(16)

Example 20-6 Creating a View Using a Hierarchical Data Guide Obtained With JSON_DATAGUIDE

This example creates a view that projects all of the fields present in the hierarchical data guide that is obtained by invoking SQL function `json_dataguide` on `po_document` of table `j_purchaseorder`. The second and third arguments passed to `json_dataguide` are used, respectively, to specify that the data guide is to be hierarchical and pretty-printed.

Note that none of the view column names here have the prefix `PO_DOCUMENT$`. That prefix is used only when you use a data guide that is obtained from the information in a data guide-enabled JSON search index.

The data guide returned by `json_dataguide` has only `o:path`, `type`, `o:length`, and `o:preferred_column_name` as its predefined fields. The values of field `o:preferred_column_name` are the same as the field names. (There can of course be no fields holding statistical information, as that can only be stored in an index.)

The view column names come from the values of field `o:preferred_column_name` of the data guide that you pass to `DBMS_JSON.create_view`. By default, the view columns are thus named the same as the projected fields.

Because the columns must be uniquely named in the view, *you must ensure* that the field names themselves are unique, or else you must edit the data guide returned by `json_dataguide` to add appropriate `o:preferred_column_name` entries that ensure uniqueness. An error is raised by `DBMS_JSON.create_view` if the names for the columns are not unique.

```

DECLARE
  dg CLOB;
BEGIN
  SELECT json_dataguide(po_document,
                       FORMAT DBMS_JSON.FORMAT_HIERARCHICAL,
                       DBMS_JSON.PRETTY)
  INTO dg
  FROM j_purchaseorder
  WHERE extract(YEAR FROM date_loaded) = 2014;
  DBMS_JSON.create_view('MYVIEW', 'J_PURCHASEORDER', 'PO_DOCUMENT',
  dg);

```

```

END;
/

DESCRIBE myview
Name                               Null?    Type
-----
DATE_LOADED                        TIMESTAMP(6) WITH TIME ZONE
ID                                  NOT NULL RAW(16)
User                                VARCHAR2(8)
PONumber                            NUMBER
UPCCode                             NUMBER
UnitPrice                           NUMBER
Description                          VARCHAR2(32)
Quantity                             NUMBER
ItemNumber                           NUMBER
Reference                            VARCHAR2(16)
Requestor                            VARCHAR2(16)
CostCenter                           VARCHAR2(4)
AllowPartialShipment                 VARCHAR2(4)
name                                  VARCHAR2(16)
Phone                                  VARCHAR2(16)
type                                  VARCHAR2(8)
number                                VARCHAR2(16)
city                                  VARCHAR2(32)
state                                  VARCHAR2(2)
street                                VARCHAR2(32)
country                               VARCHAR2(32)
zipCode                               NUMBER
Special Instructions                   VARCHAR2(8)

```

20.7.2 Creating a View Over JSON Data Based on a Path Expression

You can use the information in a data guide-enabled JSON search index to create a database view whose columns project JSON fields from your documents. The fields projected are the scalar fields not under an array plus the scalar fields in the data targeted by a specified SQL/JSON path expression.

For example, if the path expression is `$` then all scalar fields are projected, because the root (top) of the document is targeted. [Example 20-7](#) illustrates this. If the path is `$.LineItems.Part` then only the scalar fields that are present (at any level) in the data targeted by `$.LineItems.Part` are projected (in addition to scalar fields elsewhere that are not under an array). [Example 20-8](#) illustrates this.

If you gather statistics on your JSON document set then the data-guide information in a data guide-enabled JSON search index records the frequency of occurrence, across the document set, of each path to a field that is present in a document. When you create the view, you can specify that only the (scalar) fields with a given minimum frequency of occurrence (as a percentage) are to be projected as view columns. You do this by specifying a non-zero value for parameter `FREQUENCY` of procedure `DBMS_JSON.create_view_on_path`.

For example, if you specify the path as `$` and the minimum frequency as 50 then all scalar fields (on any path, since `$` targets the whole document) that occur in at least half (50%) of the documents are projected. [Example 20-9](#) illustrates this.

The value of argument *PATH* is a simple SQL/JSON path expression (no filter expression), possibly with relaxation (implicit array wrapping and unwrapping), but with no array steps and no function step. See [SQL/JSON Path Expression Syntax](#).

No frequency filtering is done in *either* of the following cases — targeted fields are *projected regardless of their frequency* of occurrence in the documents:

- You never gather statistics information on your set of JSON documents. (No frequency information is included in the data guide-enabled JSON search index.)
- The *FREQUENCY* argument of `DBMS_JSON.create_view_on_path` is zero (0).

 **Note:**

When the *FREQUENCY* argument is non-zero, even if you have gathered statistics information on your document set, the index contains *no* statistical information for any documents added after the most recent gathering of statistics. This means that any *fields added after that statistics gathering are ignored* (not projected).

 **See Also:**

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view_on_path`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`

Example 20-7 Creating a View That Projects All Scalar Fields

All scalar fields are represented in the view, because the specified path is `$`.

(Columns whose names are *italic* in the `describe` command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`. Underlined rows are missing from [Example 20-9](#).)

```
EXEC DBMS_JSON.create_view_on_path('VIEW2',
                                'J_PURCHASEORDER',
                                'PO_DOCUMENT',
                                '$');

DESCRIBE view2;
Name                               Null?    Type
-----
ID                                  NOT NULL RAW(16)
DATE_LOADED                         TIMESTAMP(6) WITH TIME ZONE
PO_DOCUMENT$User                     VARCHAR2(8)
PONumber                             NUMBER
PO_DOCUMENT$Reference                 VARCHAR2(16)
PO_DOCUMENT$Requestor                 VARCHAR2(16)
PO_DOCUMENT$CostCenter                 VARCHAR2(4)
PO_DOCUMENT$AllowPartialShipment    VARCHAR2(4)
```

PO_DOCUMENT\$name	VARCHAR2(16)
<i>Phone</i>	VARCHAR2(16)
PO_DOCUMENT\$city	VARCHAR2(32)
PO_DOCUMENT\$state	VARCHAR2(2)
PO_DOCUMENT\$street	VARCHAR2(32)
PO_DOCUMENT\$country	VARCHAR2(32)
PO_DOCUMENT\$zipCode	NUMBER
PO_DOCUMENT\$SpecialInstructions	VARCHAR2(8)
PO_DOCUMENT\$UPCCode	NUMBER
PO_DOCUMENT\$UnitPrice	NUMBER
<i>PartDescription</i>	VARCHAR2(32)
PO_DOCUMENT\$Quantity	NUMBER
<i>ItemNumber</i>	NUMBER
<i>PhoneType</i>	VARCHAR2(8)
<i>PhoneNumber</i>	VARCHAR2(16)

Example 20-8 Creating a View That Projects Scalar Fields Targeted By a Path Expression

Fields *Itemnumber*, *PhoneType*, and *PhoneNumber* are *not* represented in the view. The only fields that are projected are those scalar fields that are not under an array plus those that are present (at any level) in the data that is targeted by `$.LineItems.Part` (that is, the scalar fields whose paths start with `$.LineItems.Part`). (Columns whose names are *italic* in the describe command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`.)

```
SQL> EXEC DBMS_JSON.create_view_on_path('VIEW4',
                                         'J_PURCHASEORDER',
                                         'PO_DOCUMENT',
                                         '$.LineItems.Part');
```

```
SQL> DESCRIBE view4;
Name                                         Null?    Type
-----
ID                                           NOT NULL RAW(16)
DATE_LOADED                                TIMESTAMP(6) WITH TIME ZONE
PO_DOCUMENT$User                           VARCHAR2(8)
PONumber                                   NUMBER
PO_DOCUMENT$Reference                       VARCHAR2(16)
PO_DOCUMENT$Requestor                       VARCHAR2(16)
PO_DOCUMENT$CostCenter                      VARCHAR2(4)
PO_DOCUMENT$AllowPartialShipment            VARCHAR2(4)
PO_DOCUMENT$name                            VARCHAR2(16)
Phone                                       VARCHAR2(16)
PO_DOCUMENT$city                            VARCHAR2(32)
PO_DOCUMENT$state                           VARCHAR2(2)
PO_DOCUMENT$street                          VARCHAR2(32)
PO_DOCUMENT$country                         VARCHAR2(32)
PO_DOCUMENT$zipCode                          NUMBER
PO_DOCUMENT$SpecialInstructions              VARCHAR2(8)
PO_DOCUMENT$UPCCode                          NUMBER
PO_DOCUMENT$UnitPrice                       NUMBER
PartDescription                             VARCHAR2(32)
```

Example 20-9 Creating a View That Projects Scalar Fields Having a Given Frequency

All scalar fields that occur in all (100%) of the documents are represented in the view. Field `AllowPartialShipment` does not occur in all of the documents, so there is no column `PO_DOCUMENT$AllowPartialShipment` in the view. Similarly for fields `Phone`, `PhoneType`, and `PhoneNumber`.

(Columns whose names are *italic* in the `describe` command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`.)

```
SQL> EXEC DBMS_JSON.create_view_on_path('VIEW3',
                                        'J_PURCHASEORDER',
                                        'PO_DOCUMENT',
                                        '$',
                                        100);
```

```
SQL> DESCRIBE view3;
```

Name	Null?	Type
ID	NOT NULL	RAW(16)
DATE_LOADED		TIMESTAMP(6) WITH TIME ZONE
PO_DOCUMENT\$User		VARCHAR2(8)
<i>PONumber</i>		NUMBER
PO_DOCUMENT\$Reference		VARCHAR2(16)
PO_DOCUMENT\$Requestor		VARCHAR2(16)
PO_DOCUMENT\$CostCenter		VARCHAR2(4)
PO_DOCUMENT\$name		VARCHAR2(16)
PO_DOCUMENT\$city		VARCHAR2(32)
PO_DOCUMENT\$state		VARCHAR2(2)
PO_DOCUMENT\$street		VARCHAR2(32)
PO_DOCUMENT\$country		VARCHAR2(32)
PO_DOCUMENT\$zipCode		NUMBER
PO_DOCUMENT\$SpecialInstructions		VARCHAR2(8)
PO_DOCUMENT\$UPCCCode		NUMBER
PO_DOCUMENT\$UnitPrice		NUMBER
<i>PartDescription</i>		VARCHAR2(32)
PO_DOCUMENT\$Quantity		NUMBER
<i>ItemNumber</i>		NUMBER

Related Topics

- [Specifying a Preferred Name for a Field Column](#)
You can project JSON fields from your data as columns in a database view or as virtual columns added to the same table that contains the JSON column. You can specify a preferred name for such a column.
- [SQL/JSON Path Expressions](#)
Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.

20.8 Adding and Dropping Virtual Columns For JSON Fields Based on Data-Guide Information

Based on data-guide information for a JSON column, you can project scalar fields from that JSON data as virtual columns in the same table. The scalar fields projected are those that are not under an array.

You can do all of the following with a virtual column, with the aim of improving performance:

- Build an index on it.
- Gather statistics on it for the optimizer.
- Load it into the In-Memory Column Store (IM column store).

Note:

Virtual columns, like columns in general, are subject to the 1000-column limit for a given table.

You use PL/SQL procedure `DBMS_JSON.add_virtual_columns` to add virtual columns based on data-guide information for a JSON column. Before it adds virtual columns, procedure `add_virtual_columns` first drops any existing virtual columns that were projected from fields in the same JSON column by a previous invocation of `add_virtual_columns` or by data-guide change-trigger procedure `add_vc` (in effect, it does what procedure `DBMS_JSON.drop_virtual_columns` does).

There are two alternative sources of the data-guide information that you provide to procedure `add_virtual_columns`:

- It can come from a *hierarchical data guide* that you pass as an argument. All scalar fields in the data guide that are not under an array are projected as virtual columns. All other fields in the data guide are ignored (not projected).

In this case, you can edit the data guide before passing it, so that it specifies the scalar fields (not under an array) that you want projected. You do *not* need a data guide-enabled search index in this case.

- It can come from a *data guide-enabled JSON search index*.

In this case, you can specify, as the value of argument `FREQUENCY` to procedure `add_virtual_columns`, a minimum frequency of occurrence for the scalar fields to be projected. You need a data guide-enabled search index in this case, but you do not need a data guide.

You can also specify that added virtual columns be *hidden*. The SQL `describe` command does not list hidden columns.

- If you pass a (hierarchical) data guide to `add_virtual_columns` then you can specify projection of particular scalar fields (not under an array) as *hidden* virtual columns by adding `"o:hidden": true` to their descriptions in the data guide.

- If you use a data guide-enabled JSON search index with `add_virtual_columns` then you can specify a PL/SQL `TRUE` value for argument `HIDDEN`, to make *all* of the added virtual columns be hidden. (The default value of `HIDDEN` is `FALSE`, meaning that the added virtual columns are not hidden.)
- [Adding Virtual Columns For JSON Fields Based on a Hierarchical Data Guide](#)
You can use a hierarchical data guide to project scalar fields from JSON data as virtual columns in the same table. All scalar fields in the data guide that are not under an array are projected as virtual columns. All other fields in the data guide are ignored (not projected).
- [Adding Virtual Columns For JSON Fields Based on a Data Guide-Enabled Search Index](#)
You can use a data guide-enabled search index for a JSON column to project scalar fields from that JSON data as virtual columns in the same table. Only scalar fields not under an array are projected. You can specify a minimum frequency of occurrence for the fields to be projected.
- [Dropping Virtual Columns for JSON Fields Based on Data-Guide Information](#)
You can use procedure `DBMS_JSON.drop_virtual_columns` to drop all virtual columns that were added for JSON fields in a column of JSON data.

Related Topics

- [In-Memory JSON Data](#)
A column of JSON data can be stored in the In-Memory Column Store (IM column store) to improve query performance.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.add_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.create_view_on_path`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.drop_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`

20.8.1 Adding Virtual Columns For JSON Fields Based on a Hierarchical Data Guide

You can use a hierarchical data guide to project scalar fields from JSON data as virtual columns in the same table. All scalar fields in the data guide that are not under an array are projected as virtual columns. All other fields in the data guide are ignored (not projected).

You can obtain a hierarchical data guide using PL/SQL function `DBMS_JSON.get_index_dataguide`. A data guide-enabled JSON search index must be defined on the column of JSON data.

You can edit the data guide obtained, to include only specific scalar fields (that are not under an array), rename those fields, or change the lengths of their types. The resulting data guide specifies which such fields to project as new virtual columns. Any fields in the data guide that are not scalar fields not under an array are ignored (not projected).

You use PL/SQL procedure `DBMS_JSON.add_virtual_columns` to add the virtual columns to the table that contains the JSON column containing the projected fields. That procedure first drops any existing virtual columns that were projected from fields in the same JSON column by a previous invocation of `add_virtual_columns` or by data-guide change-trigger procedure `add_vc` (in effect, it does what procedure `DBMS_JSON.drop_virtual_columns` does).

Example 20-10 illustrates this. It projects scalar fields that are not under an array, from the data in JSON column `po_document` of table `j_purchaseorder`. The fields projected are those that are indicated in the hierarchical data guide.

Example 20-11 illustrates passing a data-guide argument that specifies the projection of two fields as virtual columns. Data-guide field `o:hidden` is used to hide one of these columns.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.add_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.drop_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`
- *Oracle Database SQL Language Reference* for information about PL/SQL constant `DBMS_JSON.FORMAT_HIERARCHICAL`

Example 20-10 Adding Virtual Columns That Project JSON Fields Using a Data Guide Obtained With GET_INDEX_DATAGUIDE

In this example the hierarchical data guide is obtained from a data guide-enabled JSON search index on JSON column `po_document`.

The added virtual columns are all of the columns in table `j_purchaseorder` except for `ID`, `DATE_LOADED`, and `PODOCUMENT`.

(Columns whose names are *italic* in the `describe` command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`.)

```
EXEC DBMS_JSON.add_virtual_columns(
    'J_PURCHASEORDER',
    'PO_DOCUMENT',
    DBMS_JSON.get_index_dataguide('J_PURCHASEORDER',
    'PO_DOCUMENT',
```

```
DBMS_JSON.FORMAT_HIERARCHICAL));
```

```
DESCRIBE j_purchaseorder;
Name                               Null?    Type
-----
ID                                  NOT NULL RAW(16)
DATE_LOADED                         TIMESTAMP(6) WITH
TIME_ZONE
PO_DOCUMENT                          CLOB
PO_DOCUMENT$User                     VARCHAR2(8)
PO_Number                             NUMBER
PO_DOCUMENT$Reference                 VARCHAR2(16)
PO_DOCUMENT$Requestor                 VARCHAR2(16)
PO_DOCUMENT$CostCenter                 VARCHAR2(4)
PO_DOCUMENT$AllowPartialShipment      VARCHAR2(4)
PO_DOCUMENT$name                       VARCHAR2(16)
Phone                                 VARCHAR2(16)
PO_DOCUMENT$city                       VARCHAR2(32)
PO_DOCUMENT$state                      VARCHAR2(2)
PO_DOCUMENT$street                     VARCHAR2(32)
PO_DOCUMENT$country                    VARCHAR2(32)
PO_DOCUMENT$zipCode                    NUMBER
PO_DOCUMENT$SpecialInstructions        VARCHAR2(8)
```

Example 20-11 Adding Virtual Columns, Hidden and Visible

In this example only two fields are projected as virtual columns: `PO_Number` and `PO_Reference`. The data guide is defined locally as a literal string. Data-guide field `o:hidden` is used here to hide the virtual column for `PO_Reference`. (For `PO_Number` the `o:hidden: false` entry is not needed, as `false` is the default value.)

```
DECLARE
  dg CLOB;
BEGIN
  dg := '{"type" : "object",
        "properties" :
          {"PO_Number" : {"type" : "number",
                          "o:length" : 4,
                          "o:preferred_column_name" : "PO_Number",
                          "o:hidden" : false},
          "PO_Reference" : {"type" : "string",
                            "o:length" : 16,
                            "o:preferred_column_name" :
"PO_Reference",
                            "o:hidden" : true}}}' ;
  DBMS_JSON.add_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT', dg);
END;
/

DESCRIBE j_purchaseorder;
Name          Null?    Type
-----
ID            NOT NULL RAW(16)
DATE_LOADED   TIMESTAMP(6) WITH TIME_ZONE
```

```

PO_DOCUMENT          CLOB
PO_Number            NUMBER

SELECT column_name FROM user_tab_columns
   WHERE table_name = 'J_PURCHASEORDER' ORDER BY 1;
COLUMN_NAME
-----
DATE_LOADED
ID
PO_DOCUMENT
PO_Number
PO_Reference

5 rows selected.

```

20.8.2 Adding Virtual Columns For JSON Fields Based on a Data Guide-Enabled Search Index

You can use a data guide-enabled search index for a JSON column to project scalar fields from that JSON data as virtual columns in the same table. Only scalar fields not under an array are projected. You can specify a minimum frequency of occurrence for the fields to be projected.

You use procedure `DBMS_JSON.add_virtual_columns` to add the virtual columns.

[Example 20-12](#) illustrates this. It projects all scalar fields that are not under an array to table `j_purchaseorder` as virtual columns.

If you gather statistics on the documents in the JSON column where you want to project fields then the data-guide information in the data guide-enabled JSON search index records the frequency of occurrence, across that document set, of each field in a document.

When you add virtual columns you can specify that only those fields with a given minimum frequency of occurrence are to be projected.

You do this by specifying a non-zero value for parameter `FREQUENCY` of procedure `add_virtual_columns`. Zero is the default value, so if you do not include argument `FREQUENCY` then all scalar fields (not under an array) are projected. The frequency of a given field is the number of documents containing that field divided by the total number of documents in the JSON column, expressed as a percentage.

[Example 20-13](#) projects all scalars (not under an array) that occur in all (100%) of the documents as virtual columns.

If you want to *hide* all of the added virtual columns then specify a `TRUE` value for argument `HIDDEN`. (The default value of parameter `HIDDEN` is `FALSE`, meaning that the added virtual columns are not hidden.)

[Example 20-14](#) projects, as hidden virtual columns, the scalar fields (not under an array) that occur in all (100%) of the documents.

 **See Also:**

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.add_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`

Example 20-12 Projecting All Scalar Fields Not Under an Array as Virtual Columns

The added virtual columns are all of the columns in table `j_purchaseorder` except for `ID`, `DATE_LOADED`, and `PODOCUMENT`. This is because no `FREQUENCY` argument is passed to `add_virtual_columns`, so all scalar fields (that are not under an array) are projected.

(Columns whose names are *italic* in the `describe` command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`.)

```
EXEC DBMS_JSON.add_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT');
```

```
DESCRIBE j_purchaseorder;
```

Name	Null?	Type
-----	-----	-----
ID	NOT NULL	RAW(16)
DATE_LOADED		TIMESTAMP(6) WITH
TIME_ZONE		
PO_DOCUMENT		CLOB
PO_DOCUMENT\$User		VARCHAR2(8)
<i>PONumber</i>		NUMBER
PO_DOCUMENT\$Reference		VARCHAR2(16)
PO_DOCUMENT\$Requestor		VARCHAR2(16)
PO_DOCUMENT\$CostCenter		VARCHAR2(4)
PO_DOCUMENT\$AllowPartialShipment		VARCHAR2(4)
PO_DOCUMENT\$name		VARCHAR2(16)
<i>Phone</i>		VARCHAR2(16)
PO_DOCUMENT\$city		VARCHAR2(32)
PO_DOCUMENT\$state		VARCHAR2(2)
PO_DOCUMENT\$street		VARCHAR2(32)
PO_DOCUMENT\$country		VARCHAR2(32)
PO_DOCUMENT\$zipCode		NUMBER
PO_DOCUMENT\$SpecialInstructions		VARCHAR2(8)

Example 20-13 Projecting Scalar Fields With a Minimum Frequency as Virtual Columns

All scalar fields that occur in all (100%) of the documents are projected as virtual columns. The result is the same as that for [Example 20-12](#), except that fields `AllowPartialShipment` and `Phone` are not projected, because they do not occur in 100% of the documents.

(Columns whose names are *italic* in the describe command output are those that have been renamed using PL/SQL procedure `DBMS_JSON.rename_column`.)

```
EXEC DBMS_JSON.add_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT',
100);

DESCRIBE j_purchaseorder;
Name                                                    Null?    Type
-----
ID                                                       NOT NULL RAW(16)
DATE_LOADED                                             TIMESTAMP(6) WITH
TIME_ZONE
PO_DOCUMENT                                             CLOB
PO_DOCUMENT$User                                       VARCHAR2(8)
PO_DOCUMENT$PONumber                                  NUMBER
PO_DOCUMENT$Reference                                 VARCHAR2(16)
PO_DOCUMENT$Requestor                                 VARCHAR2(16)
PO_DOCUMENT$CostCenter                                VARCHAR2(4)
PO_DOCUMENT$name                                       VARCHAR2(16)
PO_DOCUMENT$city                                       VARCHAR2(32)
PO_DOCUMENT$state                                       VARCHAR2(2)
PO_DOCUMENT$street                                     VARCHAR2(32)
PO_DOCUMENT$country                                   VARCHAR2(32)
PO_DOCUMENT$zipCode                                   NUMBER
PO_DOCUMENT$SpecialInstructions                       VARCHAR2(8)
```

Example 20-14 Projecting Scalar Fields With a Minimum Frequency as Hidden Virtual Columns

The result is the same as that for [Example 20-13](#), except that all of the added virtual columns are *hidden*. (The query of view `USER_TAB_COLUMNS` shows that the virtual columns were in fact added.)

```
EXEC DBMS_JSON.add_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT',
100, TRUE);

DESCRIBE j_purchaseorder;
Name                                                    Null?    Type
-----
ID                                                       NOT NULL RAW(16)
DATE_LOADED                                             TIMESTAMP(6) WITH
TIME_ZONE
PO_DOCUMENT                                             CLOB

SELECT column_name FROM user_tab_columns
  WHERE table_name = 'J_PURCHASEORDER'
     ORDER BY 1;

COLUMN_NAME
-----
DATE_LOADED
ID
```

```
PONumber  
PO_DOCUMENT  
PO_DOCUMENT$CostCenter  
PO_DOCUMENT$Reference  
PO_DOCUMENT$Requestor  
PO_DOCUMENT$SpecialInstructions  
PO_DOCUMENT$User  
PO_DOCUMENT$city  
PO_DOCUMENT$country  
PO_DOCUMENT$name  
PO_DOCUMENT$state  
PO_DOCUMENT$street  
PO_DOCUMENT$zipCode
```

20.8.3 Dropping Virtual Columns for JSON Fields Based on Data-Guide Information

You can use procedure `DBMS_JSON.drop_virtual_columns` to drop all virtual columns that were added for JSON fields in a column of JSON data.

Procedure `DBMS_JSON.drop_virtual_columns` drops all virtual columns that were projected from fields in a given JSON column by an invocation of `add_virtual_columns` or by data-guide change-trigger procedure `add_vc`.

[Example 20-15](#) illustrates this for fields projected from column `po_document` of table `j_purchaseorder`.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.add_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.drop_virtual_columns`

Example 20-15 Dropping Virtual Columns Projected From JSON Fields

```
EXEC DBMS_JSON.drop_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT');
```

20.9 Change Triggers For Data Guide-Enabled Search Index

When JSON data changes, some information in a data guide-enabled JSON search index is automatically updated. You can specify a procedure whose invocation is triggered whenever this happens. You can define your own PL/SQL procedure for this, or you can use the predefined change-trigger procedure `add_vc`.

The data-guide information in a data guide-enabled JSON search index records structure, type, and possibly statistical information about a set of JSON documents. Except for the statistical information, which is updated only when you gather statistics,

relevant changes in the document set are automatically reflected in the data-guide information stored in the index.

You can define a PL/SQL procedure whose invocation is automatically triggered by such an index update. The invocation occurs when the index is updated. Any errors that occur during the execution of the procedure are ignored.

You can use the predefined change-trigger procedure `add_vc` to automatically add virtual columns that project JSON fields from the document set or to modify existing such columns, as needed. The virtual columns added by `add_vc` follow the same naming rules as those you add by invoking procedure `DBMS_JSON.add_virtual_columns` for a JSON column that has a data guide-enabled search index.

In either case, any error that occurs during the execution of the procedure is *ignored*.

Unlike `DBMS_JSON.add_virtual_columns`, `add_vc` does *not* first drop any existing virtual columns that were projected from fields in the same JSON column. To drop virtual columns projected from fields in the same JSON column by `add_vc` or by `add_virtual_columns`, use procedure `DBMS_JSON.drop_virtual_columns`.

You specify the use of a trigger for data-guide changes by using the keywords **DATAGUIDE ON CHANGE** in the `PARAMETERS` clause when you create or alter a JSON search index. Only one change trigger is allowed per index: altering an index to specify a trigger automatically replaces any previous trigger for it.

[Example 20-16](#) alters existing JSON search index `po_search_idx` to use procedure `add_vc`.

Example 20-16 Adding Virtual Columns Automatically With Change Trigger ADD_VC

This example adds predefined change trigger `add_vc` to JSON search index `po_search_idx`.

It first drops any existing virtual columns that were projected from fields in JSON column `po_document` either by procedure `DBMS_JSON.add_virtual_columns` or by a pre-existing `add_vc` change trigger for the same JSON search index.

Then it alters the search index to add change trigger `add_vc` (if it was already present then this has no effect).

Finally, it inserts a new document that provokes a change in the data guide. Two virtual columns are added to the table, for the two scalar fields not under an array.

```
EXEC DBMS_JSON.drop_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT');

ALTER INDEX po_search_idx REBUILD PARAMETERS ('DATAGUIDE ON CHANGE
add_vc');

INSERT INTO j_purchaseorder
VALUES (
  SYS_GUID(),
  to_date('30-JUN-2015'),
  '{"PO_Number"      : 4230,
   "PO_Reference"   : "JDEER-20140421",
   "PO_LineItems"  : [{"Part_Number" : 230912362345,
                       "Quantity"   : 3.0}]}');
```

```
DESCRIBE j_purchaseorder;
Name                               Null?      Type
-----
ID                                  NOT NULL  RAW(16)
DATE_LOADED                                   TIMESTAMP(6) WITH TIME ZONE
PO_DOCUMENT                                    CLOB
PO_DOCUMENT$PO_Number                NUMBER
PO_DOCUMENT$PO_Reference              VARCHAR2(16)
```

- [User-Defined Data-Guide Change Triggers](#)
You can define a procedure whose invocation is triggered automatically whenever a given data guide-enabled JSON search index is updated. Any errors that occur during the execution of the procedure are ignored.

Related Topics

- [Adding and Dropping Virtual Columns For JSON Fields Based on Data-Guide Information](#)
Based on data-guide information for a JSON column, you can project scalar fields from that JSON data as virtual columns in the same table. The scalar fields projected are those that are not under an array.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.add_virtual_columns`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.drop_virtual_columns`

20.9.1 User-Defined Data-Guide Change Triggers

You can define a procedure whose invocation is triggered automatically whenever a given data guide-enabled JSON search index is updated. Any errors that occur during the execution of the procedure are ignored.

[Example 20-17](#) illustrates this.

A user-defined procedure specified with keywords `DATAGUIDE ON CHANGE` in a JSON search index `PARAMETERS` clause must accept the parameters specified in [Table 20-5](#).

Table 20-5 Parameters of a User-Defined Data-Guide Change Trigger Procedure

Name	Type	Description
<code>table_name</code>	VARCHAR2	Name of the table containing column <code>column_name</code> .
<code>column_name</code>	VARCHAR2	Name of a JSON column that has a data guide-enabled JSON search index.

Table 20-5 (Cont.) Parameters of a User-Defined Data-Guide Change Trigger Procedure

Name	Type	Description
path	VARCHAR2	A SQL/JSON path expression that targets a particular field in the data in column <code>column_name</code> . This path is affected by the index change that triggered the procedure invocation. For example, the index change involved adding this path or changing its type value or its type-length value.
new_type	NUMBER	A new type for the given path.
new_type_length	NUMBER	A new type length for the given path.

Example 20-17 Tracing Data-Guide Updates With a User-Defined Change Trigger

This example first drops any existing virtual columns projected from fields in JSON column `po_document`.

It then defines PL/SQL procedure `my_dataguide_trace`, which prints the names of the table and JSON column, together with the path, type and length fields of the added virtual column. It then alters JSON search index `po_search_idx` to specify that this procedure be invoked as a change trigger for updates to the data-guide information in the index.

It then inserts a new document that provokes a change in the data guide, which triggers the output of trace information.

Note that the `TYPE` argument to the procedure must be a number that is one of the `DBMS_JSON` constants for a JSON type. The procedure tests the argument and outputs a user-friendly string in place of the number.

```
EXEC DBMS_JSON.drop_virtual_columns('J_PURCHASEORDER', 'PO_DOCUMENT');

CREATE OR REPLACE PROCEDURE my_dataguide_trace(tableName VARCHAR2,
                                              jcolName VARCHAR2,
                                              path VARCHAR2,
                                              type NUMBER,
                                              tlength NUMBER)
IS
  typename VARCHAR2(10);
BEGIN
  IF (type = DBMS_JSON.TYPE_NULL) THEN typename := 'null';
  ELSIF (type = DBMS_JSON.TYPE_BOOLEAN) THEN typename := 'boolean';
  ELSIF (type = DBMS_JSON.TYPE_NUMBER) THEN typename := 'number';
  ELSIF (type = DBMS_JSON.TYPE_STRING) THEN typename := 'string';
  ELSIF (type = DBMS_JSON.TYPE_OBJECT) THEN typename := 'object';
  ELSIF (type = DBMS_JSON.TYPE_ARRAY) THEN typename := 'array';
  ELSE
    typename := 'unknown';
  END IF;
END IF;
```

```

DBMS_OUTPUT.put_line('Updating ' || tableName || '(' || jcolName
                    || '): Path = ' || path || ', Type = ' || type
                    || ', Type Name = ' || typename
                    || ', Type Length = ' || tlength);

END;
/

ALTER INDEX po_search_idx REBUILD
  PARAMETERS ('DATAGUIDE ON CHANGE my_dataguide_trace');

INSERT INTO j_purchaseorder
VALUES (
  SYS_GUID(),
  to_date('30-MAR-2016'),
  '{"PO_ID"      : 4230,
   "PO_Ref"     : "JDEER-20140421",
   "PO_Items"  : [{"Part_No"      : 98981327234,
                   "Item_Quantity" : 13}]}');

COMMIT;
Updating J_PURCHASEORDER(PO_DOCUMENT):
  Path = $.PO_ID, Type = 3, Type Name = number, Type Length = 4
Updating J_PURCHASEORDER(PO_DOCUMENT):
  Path = $.PO_Ref, Type = 4, Type Name = string, Type Length = 16
Updating J_PURCHASEORDER(PO_DOCUMENT):
  Path = $.PO_Items, Type = 6, Type Name = array, Type Length = 64
Updating J_PURCHASEORDER(PO_DOCUMENT):
  Path = $.PO_Items.Part_No, Type = 3, Type Name = number, Type Length
= 16
Updating J_PURCHASEORDER(PO_DOCUMENT):
  Path = $.PO_Items.Item_Quantity, Type = 3, Type Name = number, Type
Length = 2

Commit complete.

```

See Also:

- *Oracle Database SQL Language Reference* for information about PL/SQL constants `TYPE_NULL`, `TYPE_BOOLEAN`, `TYPE_NUMBER`, `TYPE_STRING`, `TYPE_OBJECT`, and `TYPE_ARRAY`.
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.drop_virtual_columns`

20.10 Multiple Data Guides Per Document Set

A data guide reflects the shape of a given set of JSON documents. If a JSON column contains different types of documents, with different structure or type information, you can create and use different data guides for the different kinds of documents.

Data Guides For Different Kinds of JSON Documents

JSON documents need not, and typically do not, follow a prescribed schema. This is true even for documents that are used similarly in a given application; they may differ in structural ways (shape), and field types may differ.

A JSON data guide summarizes the structural and type information of a given set of documents. In general, the more similar the structure and type information of the documents in a given set, the more useful the resulting data guide.

A data guide is created for a given column of JSON data. If the column contains very different kinds of documents (for example, purchase orders and health records) then a single data guide for the column is likely to be of limited use.

One way to address this concern is to put different kinds of JSON documents in different JSON columns. But sometimes other considerations decide in favor of mixing document types in the same column.

In addition, documents of the same general type, which you decide to store in the same column, can nevertheless differ in relatively systematic ways. This includes the case of *evolving* document shape and type information. For example, the structure of tax-information documents could change from year to year.

When you create a data guide you can decide which information to summarize. And you can thus create different data guides for the same JSON column, to represent different subsets of the document set.

An additional aid in this regard is to have a separate, non-JSON, column in the same table, which is used to label, or categorize, the documents in a JSON column.

In the case of the purchase-order documents used in our examples, let's suppose that their structure can evolve significantly from year to year, so that column `date_loaded` of table `j_purchaseorder` can be used to group them into subsets of reasonably similar shape. [Example 20-18](#) adds a purchase-order document for 2015, and [Example 20-19](#) adds a purchase-order document for 2016. (Compare with the documents for 2014, which are added in [Example 4-2](#).)

Using a SQL Aggregate Function to Create Multiple Data Guides

Oracle SQL function `json_dataguide` is in fact an *aggregate* function. An aggregate function returns a single result row based on groups of rows, rather than on a single row. It is typically used in a `SELECT` list for a query that has a `GROUP BY` clause, which divides the rows of a queried table or view into groups. The aggregate function applies to each group of rows, returning a single result row for each group. For example, aggregate function `avg` returns the average of a group of values.

Function `json_dataguide` aggregates JSON data to produce a summary, or specification, of it, which is returned in the form of a JSON document. In other words, for each group of JSON documents to which they are applied, they return a data guide.

If you omit `GROUP BY` then this function returns a single data guide that summarizes all of the JSON data in the subject JSON column.

Example 20-20 queries the documents of JSON column `po_document`, grouping them to produce three data guides, one for each year of column `date_loaded`.

Example 20-18 Adding a 2015 Purchase-Order Document

The 2015 purchase-order format uses only part number, reference, and line-items as its top-level fields, and these fields use prefix `PO_`. Each line item contains only a part number and a quantity.

```
INSERT INTO j_purchaseorder
VALUES (
  SYS_GUID(),
  to_date('30-JUN-2015'),
  '{"PO_Number"      : 4230,
   "PO_Reference"    : "JDEER-20140421",
   "PO_LineItems"   : [{"Part_Number" : 230912362345,
                        "Quantity"   : 3.0}]}');
```

Example 20-19 Adding a 2016 Purchase-Order Document

The 2016 format uses `PO_ID` instead of `PO_Number`, `PO_Ref` instead of `PO_Reference`, `PO_Items` instead of `PO_LineItems`, `Part_No` instead of `Part_Number`, and `Item_Quantity` instead of `Quantity`.

```
INSERT INTO j_purchaseorder
VALUES (
  SYS_GUID(),
  to_date('30-MAR-2016'),
  '{"PO_ID"         : 4230,
   "PO_Ref"        : "JDEER-20140421",
   "PO_Items"     : [{"Part_No"       : 98981327234,
                      "Item_Quantity" : 13}]}');
```

Example 20-20 Creating Multiple Data Guides With Aggregate Function JSON_DATAGUIDE

This example uses aggregate SQL function `json_dataguide` to obtain three flat¹ data guides, one for each year-specific format. The data guide for 2014 is shown only partially — it is the same as the data guide from [A Flat Data Guide For Purchase-Order Documents](#), except that no statistics fields are present. (Data guides returned by functions `json_dataguide` do not contain any statistics fields.)

```
SELECT extract(YEAR FROM date_loaded), json_dataguide(po_document)
FROM j_purchaseorder
  GROUP BY extract(YEAR FROM date_loaded)
  ORDER BY extract(YEAR FROM date_loaded) DESC;

EXTRACT(YEARFROMDATE_LOADED)
-----
```

¹ If function `json_dataguide` were passed `DBMS_JSON.FORMAT_HIERARCHICAL` as optional second argument then the result would be three hierarchical data guides.

JSON_DATAGUIDE(PO_DOCUMENT)

```
-----  
2016  
[  
  {  
    "o:path" : "$.PO_ID",  
    "type" : "number",  
    "o:length" : 4  
  },  
  {  
    "o:path" : "$.PO_Ref",  
    "type" : "string",  
    "o:length" : 16  
  },  
  {  
    "o:path" : "$.PO_Items",  
    "type" : "array",  
    "o:length" : 64  
  },  
  {  
    "o:path" : "$.PO_Items.Part_No",  
    "type" : "number",  
    "o:length" : 16  
  },  
  {  
    "o:path" : "$.PO_Items.Item_Quantity",  
    "type" : "number",  
    "o:length" : 2  
  }  
]  
  
2015  
[  
  {  
    "o:path" : "$.PO_Number",  
    "type" : "number",  
    "o:length" : 4  
  },  
  {  
    "o:path" : "$.PO_LineItems",  
    "type" : "array",  
    "o:length" : 64  
  },  
  {  
    "o:path" : "$.PO_LineItems.Quantity",  
    "type" : "number",  
    "o:length" : 4  
  },  
  {  
    "o:path" : "$.PO_LineItems.Part_Number",  
    "type" : "number",  
    "o:length" : 16  
  },  
  {  
    "o:path" : "$.PO_Reference",
```

```
    "type" : "string",
    "o:length" : 16
  }
]

                                2014

[
  {
    "o:path" : "$.User",
    "type" : "string",
    "o:length" : 8
  },
  {
    "o:path" : "$.PONumber",
    "type" : "number",
    "o:length" : 4
  },
  ...
  {
    "o:path" : "$.\"Special Instructions\"",
    "type" : "string",
    "o:length" : 8
  }
]
```

3 rows selected.



See Also:

Oracle Database SQL Language Reference for information about SQL function `json_dataguide`

20.11 Querying a Data Guide

A data guide is information about a set of JSON documents. You can query it from a flat data guide that you obtain using either Oracle SQL function `json_dataguide`

or PL/SQL function `DBMS_JSON.get_index_dataguide`. In the latter case, a data guide-enabled JSON search index must be defined on the JSON data.

See Also:

- *Oracle Database SQL Language Reference* for information about SQL function `json_dataguide`
- *Oracle Database SQL Language Reference* for information about SQL/JSON function `json_table`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database SQL Language Reference* for information about PL/SQL constant `DBMS_JSON.FORMAT_FLAT`

Example 20-21 Querying a Data Guide Obtained Using `JSON_DATAGUIDE`

This example uses SQL/JSON function `json_dataguide` to obtain a flat data guide. It then queries the relational columns projected on the fly by SQL/JSON function `json_table` from fields `o:path`, `type`, and `o:length`. It returns the projected columns ordered lexicographically by the path column created, `jpath`.

```
WITH dg_t AS (SELECT json_dataguide(po_document) dg_doc FROM
j_purchaseorder)
  SELECT jt.*
     FROM dg_t,
          json_table(dg_doc, '$[*]'
                     COLUMNS
                        jpath  VARCHAR2(40) PATH '$."o:path"',
                        type   VARCHAR2(10) PATH '$."type"',
                        tlength NUMBER      PATH '$."o:length"') jt
     ORDER BY jt.jpath;
```

JPATH	TYPE	TLENGTH
\$. "Special Instructions"	string	8
\$. AllowPartialShipment	boolean	4
\$. CostCenter	string	4
\$. LineItems	array	512
\$. LineItems.ItemNumber	number	1
\$. LineItems.Part	object	128
\$. LineItems.Part.Description	string	32
\$. LineItems.Part.UPCCode	number	16
\$. LineItems.Part.UnitPrice	number	8
\$. LineItems.Quantity	number	4
\$. PONumber	number	4
\$. PO_LineItems	array	64
\$. Reference	string	16
\$. Requestor	string	16
\$. ShippingInstructions	object	256
\$. ShippingInstructions.Address	object	128

\$.ShippingInstructions.Address.city	string	32
\$.ShippingInstructions.Address.country	string	32
\$.ShippingInstructions.Address.state	string	2
\$.ShippingInstructions.Address.street	string	32
\$.ShippingInstructions.Address.zipCode	number	8
\$.ShippingInstructions.Phone	array	128
\$.ShippingInstructions.Phone	string	16
\$.ShippingInstructions.Phone.number	string	16
\$.ShippingInstructions.Phone.type	string	8
\$.ShippingInstructions.name	string	16
\$.User	string	8

Example 20-22 Querying a Data Guide With Index Data For Paths With Frequency at Least 80%

This example uses PL/SQL function `DBMS_JSON.get_index_dataguide` with format value `DBMS_JSON.FORMAT_FLAT` to obtain a flat data guide from the data-guide information stored in a data guide-enabled JSON search index. It then queries the relational columns projected on the fly from fields `o:path`, `type`, `o:length`, and `o:frequency` by SQL/JSON function `json_table`.

The value of field `o:frequency` is a statistic that records the frequency of occurrence, across the document set, of each field in a document. It is available *only if you have gathered statistics* on the document set. The frequency of a given field is the number of documents containing that field divided by the total number of documents in the JSON column, expressed as a percentage.

```
WITH dg_t AS (SELECT DBMS_JSON.get_index_dataguide('J_PURCHASEORDER',
                                                'PO_DOCUMENT',
```

```
DBMS_JSON.FORMAT_FLAT) dg_doc
      FROM DUAL)
SELECT jt.*
   FROM dg_t,
      json_table(dg_doc, '$[*]'
      COLUMNS
        jpath    VARCHAR2(40) PATH '$."o:path"',
        type     VARCHAR2(10) PATH '$."type"',
        tlength  NUMBER      PATH '$."o:length"',
        frequency NUMBER      PATH '$."o:frequency"') jt
   WHERE jt.frequency > 80;
```

JPATH	TYPE	TLENGTH
FREQUENCY		

\$.User	string	8
100		
\$.PONumber	number	4
100		
\$.LineItems	array	512
100		
\$.LineItems.Part	object	128
100		
\$.LineItems.Part.UPCCode	number	16

100			
100	\$.LineItems.Part.UnitPrice	number	8
100	\$.LineItems.Part.Description	string	32
100	\$.LineItems.Quantity	number	4
100	\$.LineItems.ItemNumber	number	1
100	\$.Reference	string	16
100	\$.Requestor	string	16
100	\$.CostCenter	string	4
100	\$.ShippingInstructions	object	256
100	\$.ShippingInstructions.name	string	16
100	\$.ShippingInstructions.Address	object	128
100	\$.ShippingInstructions.Address.city	string	32
100	\$.ShippingInstructions.Address.state	string	2
100	\$.ShippingInstructions.Address.street	string	32
100	\$.ShippingInstructions.Address.country	string	32
100	\$.ShippingInstructions.Address.zipCode	number	8
100	\$. "Special Instructions"	string	8
100			

Related Topics

- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.

20.12 A Flat Data Guide For Purchase-Order Documents

The fields of a sample flat data guide are described. It corresponds to a set of purchase-order documents.

The only JSON Schema keyword used in a flat data guide is **type**. The other fields are all Oracle data-guide fields, which have prefix **o:**.

[Example 20-23](#) shows a flat data guide for the purchase-order documents in table `j_purchaseorder`. Things to note:

- The values of `o:preferred_column_name` use prefix `PO_DOCUMENT$`. This prefix comes from using `DBMS_JSON.get_index_dataguide` to obtain this data guide.

- The value of `o:length` is 8 for path `$.User`, for example, in spite of the fact that the actual lengths of the field values are 5. This is because the value of `o:length` is always a power of two.
- The value of `o:path` for field `Special Instructions` is wrapped in double quotation marks ("`Special Instructions`") because of the embedded space character.

Example 20-23 Flat Data Guide For Purchase Orders

Paths are **bold**. JSON schema keywords are *italic*. Preferred column names that result from using `DBMS_JSON.rename_column` are also *italic*. The formatting used is similar to that produced by using `SQL/JSON` function `json_dataguide` with format arguments `DBMS_JSON.FORMAT_FLAT` and `DBMS_JSON.PRETTY`.

Note that fields `o:frequency`, `o:low_value`, `o:high_value`, `o:num_nulls`, and `o:last_analyzed` are present. This can only be because statistics were gathered on the document set. Their values reflect the state as of the last statistics gathering. See [Example 20-3](#) for an example of gathering statistics for this data.

```
[
  {
    "o:path": "$.User",
    "type": "string",
    "o:length": 8,
    "o:preferred_column_name": "PO_DOCUMENT$User",
    "o:frequency": 100,
    "o:low_value": "ABULL",
    "o:high_value": "SBELL",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.PONumber",
    "type": "number",
    "o:length": 4,
    "o:preferred_column_name": "PONumber",
    "o:frequency": 100,
    "o:low_value": "672",
    "o:high_value": "1600",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.LineItems",
    "type": "array",
    "o:length": 512,
    "o:preferred_column_name": "PO_DOCUMENT$LineItems",
    "o:frequency": 100,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.LineItems.Part",
    "type": "object",
    "o:length": 128,
    "o:preferred_column_name": "PO_DOCUMENT$Part",

```

```

"o:frequency": 100,
"o:last_analyzed": "2016-03-31T12:17:53"
},
{
"o:path": "$.LineItems.Part.UPCCode",
"type": "number",
"o:length": 16,
"o:preferred_column_name": "PO_DOCUMENT$UPCCode",
"o:frequency": 100,
"o:low_value": "13131092899",
"o:high_value": "717951002396",
"o:num_nulls": 0,
"o:last_analyzed": "2016-03-31T12:17:53"
},
{
"o:path": "$.LineItems.Part.UnitPrice",
"type": "number",
"o:length": 8,
"o:preferred_column_name": "PO_DOCUMENT$UnitPrice",
"o:frequency": 100,
"o:low_value": "20",
"o:high_value": "19.95",
"o:num_nulls": 0,
"o:last_analyzed": "2016-03-31T12:17:53"
},
{
"o:path": "$.LineItems.Part.Description",
"type": "string",
"o:length": 32,
"o:preferred_column_name": "PartDescription",
"o:frequency": 100,
"o:low_value": "Nixon",
"o:high_value": "Eric Clapton: Best Of 1981-1999",
"o:num_nulls": 0,
"o:last_analyzed": "2016-03-31T12:17:53"
},
{
"o:path": "$.LineItems.Quantity",
"type": "number",
"o:length": 4,
"o:preferred_column_name": "PO_DOCUMENT$Quantity",
"o:frequency": 100,
"o:low_value": "5",
"o:high_value": "9.0",
"o:num_nulls": 0,
"o:last_analyzed": "2016-03-31T12:17:53"
},
{
"o:path": "$.LineItems.ItemNumber",
"type": "number",
"o:length": 1,
"o:preferred_column_name": "ItemNumber",
"o:frequency": 100,
"o:low_value": "1",
"o:high_value": "3",

```



```

    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.Reference",
    "type": "string",
    "o:length": 16,
    "o:preferred_column_name": "PO_DOCUMENT$Reference",
    "o:frequency": 100,
    "o:low_value": "ABULL-20140421",
    "o:high_value": "SBELL-20141017",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.Requestor",
    "type": "string",
    "o:length": 16,
    "o:preferred_column_name": "PO_DOCUMENT$Requestor",
    "o:frequency": 100,
    "o:low_value": "Sarah Bell",
    "o:high_value": "Alexis Bull",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.CostCenter",
    "type": "string",
    "o:length": 4,
    "o:preferred_column_name": "PO_DOCUMENT$CostCenter",
    "o:frequency": 100,
    "o:low_value": "A50",
    "o:high_value": "A50",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.AllowPartialShipment",
    "type": "boolean",
    "o:length": 4,
    "o:preferred_column_name": "PO_DOCUMENT$AllowPartialShipment",
    "o:frequency": 50,
    "o:low_value": "true",
    "o:high_value": "true",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.ShippingInstructions",
    "type": "object",
    "o:length": 256,
    "o:preferred_column_name": "PO_DOCUMENT$ShippingInstructions",
    "o:frequency": 100,
    "o:last_analyzed": "2016-03-31T12:17:53"
  }
},

```

```
{
  "o:path": "$.ShippingInstructions.name",
  "type": "string",
  "o:length": 16,
  "o:preferred_column_name": "PO_DOCUMENT$name",
  "o:frequency": 100,
  "o:low_value": "Sarah Bell",
  "o:high_value": "Alexis Bull",
  "o:num_nulls": 0,
  "o:last_analyzed": "2016-03-31T12:17:53"
},
{
  "o:path": "$.ShippingInstructions.Phone",
  "type": "string",
  "o:length": 16,
  "o:preferred_column_name": "Phone",
  "o:frequency": 50,
  "o:low_value": "983-555-6509",
  "o:high_value": "983-555-6509",
  "o:num_nulls": 0,
  "o:last_analyzed": "2016-03-31T12:17:53"
},
{
  "o:path": "$.ShippingInstructions.Phone",
  "type": "array",
  "o:length": 128,
  "o:preferred_column_name": "PO_DOCUMENT$Phone_1",
  "o:frequency": 50,
  "o:last_analyzed": "2016-03-31T12:17:53"
},
{
  "o:path": "$.ShippingInstructions.Phone.type",
  "type": "string",
  "o:length": 8,
  "o:preferred_column_name": "PhoneType",
  "o:frequency": 50,
  "o:low_value": "Mobile",
  "o:high_value": "Office",
  "o:num_nulls": 0,
  "o:last_analyzed": "2016-03-31T12:17:53"
},
{
  "o:path": "$.ShippingInstructions.Phone.number",
  "type": "string",
  "o:length": 16,
  "o:preferred_column_name": "PhoneNumber",
  "o:frequency": 50,
  "o:low_value": "415-555-1234",
  "o:high_value": "909-555-7307",
  "o:num_nulls": 0,
  "o:last_analyzed": "2016-03-31T12:17:53"
},
{
  "o:path": "$.ShippingInstructions.Address",
  "type": "object",
```

```

    "o:length": 128,
    "o:preferred_column_name": "PO_DOCUMENT$Address",
    "o:frequency": 100,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.ShippingInstructions.Address.city",
    "type": "string",
    "o:length": 32,
    "o:preferred_column_name": "PO_DOCUMENT$city",
    "o:frequency": 100,
    "o:low_value": "South San Francisco",
    "o:high_value": "South San Francisco",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.ShippingInstructions.Address.state",
    "type": "string",
    "o:length": 2,
    "o:preferred_column_name": "PO_DOCUMENT$state",
    "o:frequency": 100,
    "o:low_value": "CA",
    "o:high_value": "CA",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.ShippingInstructions.Address.street",
    "type": "string",
    "o:length": 32,
    "o:preferred_column_name": "PO_DOCUMENT$street",
    "o:frequency": 100,
    "o:low_value": "200 Sporting Green",
    "o:high_value": "200 Sporting Green",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.ShippingInstructions.Address.country",
    "type": "string",
    "o:length": 32,
    "o:preferred_column_name": "PO_DOCUMENT$country",
    "o:frequency": 100,
    "o:low_value": "United States of America",
    "o:high_value": "United States of America",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.ShippingInstructions.Address.zipCode",
    "type": "number",
    "o:length": 8,
    "o:preferred_column_name": "PO_DOCUMENT$zipCode",
    "o:frequency": 100,

```

```

    "o:low_value": "99236",
    "o:high_value": "99236",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "o:path": "$.\"Special Instructions\"",
    "type": "string",
    "o:length": 8,
    "o:preferred_column_name": "PO_DOCUMENT$SpecialInstructions",
    "o:frequency": 100,
    "o:low_value": "Courier",
    "o:high_value": "Courier",
    "o:num_nulls": 1,
    "o:last_analyzed": "2016-03-31T12:17:53"
  }
]

```

Related Topics

- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.
- [Specifying a Preferred Name for a Field Column](#)
You can project JSON fields from your data as columns in a database view or as virtual columns added to the same table that contains the JSON column. You can specify a preferred name for such a column.

See Also:

- [Example 4-2](#)
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.get_index_dataguide`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`

20.13 A Hierarchical Data Guide For Purchase-Order Documents

The fields of a sample hierarchical data guide are described. It corresponds to a set of purchase-order documents.

[Example 20-24](#) shows a hierarchical data guide for the purchase-order documents in table `j_purchaseorder`.

Example 20-24 Hierarchical Data Guide For Purchase Orders

Field names are **bold**. JSON Schema keywords are *italic*. Preferred column names that result from using `DBMS_JSON.rename_column` are also *italic*. The formatting used

is similar to that produced by using SQL/JSON function `json_dataguide` with format arguments `DBMS_JSON.FORMAT_HIERARCHICAL` and `DBMS_JSON.PRETTY`.

Note that fields `o:frequency`, `o:low_value`, `o:high_value`, `o:num_nulls`, and `o:last_analyzed` are present in this example. This can only be because statistics were gathered on the document set. Their values reflect the state as of the last statistics gathering. See [Example 20-3](#) for an example of gathering statistics for this data.

```
{
  "type": "object",
  "properties": {
    "User": {
      "type": "string",
      "o:length": 8,
      "o:preferred_column_name": "PO_DOCUMENT$User",
      "o:frequency": 100,
      "o:low_value": "ABULL",
      "o:high_value": "SBELL",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "PONumber": {
      "type": "number",
      "o:length": 4,
      "o:preferred_column_name": "PONumber",
      "o:frequency": 100,
      "o:low_value": "672",
      "o:high_value": "1600",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "LineItems": {
      "type": "array",
      "o:length": 512,
      "o:preferred_column_name": "PO_DOCUMENT$LineItems",
      "o:frequency": 100,
      "o:last_analyzed": "2016-03-31T12:17:53",
      "items": {
        "properties": {
          "Part": {
            "type": "object",
            "o:length": 128,
            "o:preferred_column_name": "PO_DOCUMENT$Part",
            "o:frequency": 100,
            "o:last_analyzed": "2016-03-31T12:17:53",
            "properties": {
              "UPCCode": {
                "type": "number",
                "o:length": 16,
                "o:preferred_column_name": "PO_DOCUMENT$UPCCode",
                "o:frequency": 100,
                "o:low_value": "13131092899",
                "o:high_value": "717951002396",
                "o:num_nulls": 0,

```

```

        "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "UnitPrice": {
        "type": "number",
        "o:length": 8,
        "o:preferred_column_name": "PO_DOCUMENT$UnitPrice",
        "o:frequency": 100,
        "o:low_value": "20",
        "o:high_value": "19.95",
        "o:num_nulls": 0,
        "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "Description": {
        "type": "string",
        "o:length": 32,
        "o:preferred_column_name": "PartDescription",
        "o:frequency": 100,
        "o:low_value": "Nixon",
        "o:high_value": "Eric Clapton: Best Of 1981-1999",
        "o:num_nulls": 0,
        "o:last_analyzed": "2016-03-31T12:17:53"
    }
}
},
"Quantity": {
    "type": "number",
    "o:length": 4,
    "o:preferred_column_name": "PO_DOCUMENT$Quantity",
    "o:frequency": 100,
    "o:low_value": "5",
    "o:high_value": "9.0",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
},
"ItemNumber": {
    "type": "number",
    "o:length": 1,
    "o:preferred_column_name": "ItemNumber",
    "o:frequency": 100,
    "o:low_value": "1",
    "o:high_value": "3",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
}
}
},
"Reference": {
    "type": "string",
    "o:length": 16,
    "o:preferred_column_name": "PO_DOCUMENT$Reference",
    "o:frequency": 100,
    "o:low_value": "ABULL-20140421",
    "o:high_value": "SBELL-20141017",
    "o:num_nulls": 0,

```

```

    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  "Requestor": {
    "type": "string",
    "o:length": 16,
    "o:preferred_column_name": "PO_DOCUMENT$Requestor",
    "o:frequency": 100,
    "o:low_value": "Sarah Bell",
    "o:high_value": "Alexis Bull",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  "CostCenter": {
    "type": "string",
    "o:length": 4,
    "o:preferred_column_name": "PO_DOCUMENT$CostCenter",
    "o:frequency": 100,
    "o:low_value": "A50",
    "o:high_value": "A50",
    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  "AllowPartialShipment": {
    "type": "boolean",
    "o:length": 4,
    "o:preferred_column_name": "PO_DOCUMENT$AllowPartialShipment",
    "o:frequency": 50,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  "ShippingInstructions": {
    "type": "object",
    "o:length": 256,
    "o:preferred_column_name": "PO_DOCUMENT$ShippingInstructions",
    "o:frequency": 100,
    "o:last_analyzed": "2016-03-31T12:17:53",
    "properties": {
      "name": {
        "type": "string",
        "o:length": 16,
        "o:preferred_column_name": "PO_DOCUMENT$name",
        "o:frequency": 100,
        "o:low_value": "Sarah Bell",
        "o:high_value": "Alexis Bull",
        "o:num_nulls": 0,
        "o:last_analyzed": "2016-03-31T12:17:53"
      },
      "Phone": {
        "oneOf": [
          {
            "type": "string",
            "o:length": 16,
            "o:preferred_column_name": "Phone",
            "o:frequency": 50,
            "o:low_value": "983-555-6509",
            "o:high_value": "983-555-6509",

```

```

    "o:num_nulls": 0,
    "o:last_analyzed": "2016-03-31T12:17:53"
  },
  {
    "type": "array",
    "o:length": 128,
    "o:preferred_column_name": "PO_DOCUMENT$Phone_1",
    "o:frequency": 50,
    "o:last_analyzed": "2016-03-31T12:17:53",
    "items": {
      "properties": {
        "type": {
          "type": "string",
          "o:length": 8,
          "o:preferred_column_name": "PhoneType",
          "o:frequency": 50,
          "o:low_value": "Mobile",
          "o:high_value": "Office",
          "o:num_nulls": 0,
          "o:last_analyzed": "2016-03-31T12:17:53"
        },
        "number": {
          "type": "string",
          "o:length": 16,
          "o:preferred_column_name": "PhoneNumber",
          "o:frequency": 50,
          "o:low_value": "415-555-1234",
          "o:high_value": "909-555-7307",
          "o:num_nulls": 0,
          "o:last_analyzed": "2016-03-31T12:17:53"
        }
      }
    }
  }
]
},
"Address": {
  "type": "object",
  "o:length": 128,
  "o:preferred_column_name": "PO_DOCUMENT$Address",
  "o:frequency": 100,
  "o:last_analyzed": "2016-03-31T12:17:53",
  "properties": {
    "city": {
      "type": "string",
      "o:length": 32,
      "o:preferred_column_name": "PO_DOCUMENT$city",
      "o:frequency": 100,
      "o:low_value": "South San Francisco",
      "o:high_value": "South San Francisco",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "state": {
      "type": "string",

```



```

      "o:length": 2,
      "o:preferred_column_name": "PO_DOCUMENT$state",
      "o:frequency": 100,
      "o:low_value": "CA",
      "o:high_value": "CA",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "street": {
      "type": "string",
      "o:length": 32,
      "o:preferred_column_name": "PO_DOCUMENT$street",
      "o:frequency": 100,
      "o:low_value": "200 Sporting Green",
      "o:high_value": "200 Sporting Green",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "country": {
      "type": "string",
      "o:length": 32,
      "o:preferred_column_name": "PO_DOCUMENT$country",
      "o:frequency": 100,
      "o:low_value": "United States of America",
      "o:high_value": "United States of America",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    },
    "zipCode": {
      "type": "number",
      "o:length": 8,
      "o:preferred_column_name": "PO_DOCUMENT$zipCode",
      "o:frequency": 100,
      "o:low_value": "99236",
      "o:high_value": "99236",
      "o:num_nulls": 0,
      "o:last_analyzed": "2016-03-31T12:17:53"
    }
  },
  "Special Instructions": {
    "type": "string",
    "o:length": 8,
    "o:preferred_column_name": "PO_DOCUMENT$SpecialInstructions",
    "o:frequency": 100,
    "o:low_value": "Courier",
    "o:high_value": "Courier",
    "o:num_nulls": 1,
    "o:last_analyzed": "2016-03-31T12:17:53"
  }
}

```

Related Topics

- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.
- [Specifying a Preferred Name for a Field Column](#)
You can project JSON fields from your data as columns in a database view or as virtual columns added to the same table that contains the JSON column. You can specify a preferred name for such a column.

 **See Also:**

- [Example 4-2](#)
- *Oracle Database SQL Language Reference* for information about SQL function `json_dataguide`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `DBMS_JSON.rename_column`

Part V

Generation of JSON Data

You can use SQL to generate JSON data from other kinds of database data programmatically, using SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg`.

- [Generation of JSON Data with SQL/JSON Functions](#)
SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` are presented.

21

Generation of JSON Data with SQL/JSON Functions

SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` are presented.

- [Overview of SQL/JSON Generation Functions](#)
You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.
- [Handling of Input Values For SQL/JSON Generation Functions](#)
The SQL/JSON generation functions take SQL values as input and render them as JSON values inside the SQL value that is returned. How the input values are rendered as JSON depends on their SQL data type.
- [SQL/JSON Function JSON_OBJECT](#)
SQL/JSON function `json_object` constructs JSON objects from the results of evaluating its argument SQL expressions.
- [SQL/JSON Function JSON_ARRAY](#)
SQL/JSON function `json_array` constructs a JSON array from the results of evaluating its argument SQL expressions.
- [SQL/JSON Function JSON_OBJECTAGG](#)
SQL/JSON function `json_objectagg` constructs a JSON object by aggregating information from multiple rows of a grouped SQL query as the object members.
- [SQL/JSON Function JSON_ARRAYAGG](#)
SQL/JSON function `json_arrayagg` constructs a JSON array by aggregating information from multiple rows of a grouped SQL query as the array elements. The order of array elements reflects the query result order, by default, but you can use the `ORDER BY` clause to impose array element order.

21.1 Overview of SQL/JSON Generation Functions

You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.

These generation functions make it easy to construct JSON data directly from a SQL query. They allow non-JSON data to be represented as JSON objects and JSON arrays. You can generate complex, hierarchical JSON documents by nesting calls to these functions. Nested subqueries can generate JSON collections that represent one-to-many relationships.¹

¹ The behavior of the SQL/JSON generation functions for JSON data is similar to that of the SQL/XML generation functions for XML data.

The Best Way to Construct JSON Data from Non-JSON Data

Alternatives to using the SQL/JSON generation functions are generally error prone or inefficient.

- Using *string concatenation* to generate JSON documents is error prone. In particular, there are a number of complex rules that must be respected concerning when and how to escape special characters, such as double quotation marks ("). It is easy to overlook or misunderstand these rules, which can result in generating incorrect JSON data.
- Reading non-JSON result sets from the database and using *client-side application code* to generate JSON data is typically quite inefficient, particularly due to network overhead. When representing one-to-many relationships as JSON data, multiple `SELECT` operations are often required, to collect all of the non-JSON data needed. If the documents to be generated represent multiple levels of one-to-many relationships then this technique can be quite costly.

The SQL/JSON generation functions do not suffer from such problems; they are designed for the job of constructing JSON data from non-JSON database data.

- They always construct well-formed JSON documents.
- By using SQL subqueries with these functions, you can generate an entire set of JSON documents using a single SQL statement, which allows the generation operation to be optimized.
- Because only the generated documents are returned to a client, network overhead is minimized: there is at most one round trip per document generated.

The SQL/JSON Generation Functions

- Functions `json_object` and `json_array` construct a JSON object or array, respectively. In the simplest case, `json_object` takes SQL name–value pairs as arguments, and `json_array` takes SQL values as arguments.
- Functions `json_objectagg`, and `json_arrayagg` are *aggregate* SQL functions. They transform information that is contained in the rows of a grouped SQL query into JSON objects and arrays, respectively. Evaluation of the arguments determines the number of object members and array elements, respectively; that is, the size of the result reflects the current queried data.

For `json_objectagg` and `json_arrayagg`, the order of object members and array elements, respectively, is unspecified. For `json_arrayagg`, you can use an `ORDER BY` clause within the `json_arrayagg` invocation to control the array element order.

Result Returned by SQL/JSON Generation Functions

By default, the generated JSON data is returned from a generation function as a SQL `VARCHAR2` value. You can use the optional `RETURNING` clause to specify a `VARCHAR2` size or to specify a `CLOB` or `BLOB` return value instead. When `BLOB` is the return type, the character set is `AL32UTF8`.

Handling of Input Values For SQL/JSON Generation Functions

The SQL/JSON generation functions take SQL values as input and render them as JSON values inside the SQL value that is returned. How the input values are rendered as JSON depends on their SQL data type.

Optional Behavior For SQL/JSON Generation Functions

You can optionally specify a SQL `NULL`-handling clause, a `RETURNING` clause, and keyword `STRICT`.

- **NULL-handling clause** — Determines how a SQL `NULL` value resulting from input evaluation is handled.
 - **NULL ON NULL** — An input SQL `NULL` value is converted to JSON `null` for output. This is the default behavior for `json_object` and `json_objectagg`.
 - **ABSENT ON NULL** — An input SQL `NULL` value results in no corresponding output. This is the default behavior for `json_array` and `json_arrayagg`.
- **RETURNING clause** — The SQL data type used for the function return value. The default is `VARCHAR2(4000)`.
- **STRICT keyword** — If present, the returned JSON data is checked, to be sure it is well-formed. If `STRICT` is present and the returned data is not well-formed then an error is raised.

Related Topics

- [Handling of Input Values For SQL/JSON Generation Functions](#)
The SQL/JSON generation functions take SQL values as input and render them as JSON values inside the SQL value that is returned. How the input values are rendered as JSON depends on their SQL data type.
- [ISO 8601 Date and Time Support](#)
International Standards Organization (ISO) standard 8601 describes an internationally accepted way to represent dates and times. Oracle Database supports many of the ISO 8601 date and time formats.

See Also:

- *Oracle Database SQL Language Reference* for information about SQL/JSON function `json_array`
- *Oracle Database SQL Language Reference* for information about SQL/JSON function `json_arrayagg`
- *Oracle Database SQL Language Reference* for information about SQL/JSON function `json_object`
- *Oracle Database SQL Language Reference* for information about SQL/JSON function `json_objectagg`

21.2 Handling of Input Values For SQL/JSON Generation Functions

The SQL/JSON generation functions take SQL values as input and render them as JSON values inside the SQL value that is returned. How the input values are rendered as JSON depends on their SQL data type.

By default, a SQL `NUMBER` value is rendered in the output as a JSON number, a SQL `VARCHAR2` value is rendered as a JSON string, and so on. For example, by default the `VARCHAR2` value `' {} '` is rendered as the JSON string `" {} "`.

In some cases you know or expect that an input value in fact already represents JSON data, and you want to communicate this to the generation function so that the value is kept as is. For example, if the input is `' {} '` then you might want it interpreted (rendered) as an empty JSON *object*, `{}`, not as a JSON *string*, `" {} "`.

You can do this by adding keywords `FORMAT JSON` after an input expression to declare that the value that results from it is to be considered as already representing JSON data. Equivalently, you can apply SQL function `treat` with keywords `AS JSON` to a generation-function input value — the effect is the same as using `FORMAT JSON`.

In many cases Oracle can determine automatically that an input value is in fact JSON data, in which case the input is treated as if it were followed by an explicit `FORMAT JSON` declaration. This is the case, for instance, if the value expression is itself an invocation of a SQL/JSON generation function or function `json_query`.

If you do *not* specify `FORMAT JSON` for a given input value, and if Oracle *cannot* determine that the value is JSON data, then it is assumed to be ordinary (non-JSON) SQL data. In that case it is serialized as follows (any other SQL value raises an error):

- An instance of a user-defined SQL object type is rendered as a JSON object whose field names are taken from the object attribute names and whose field values are taken from the object attribute values (to which JSON generation is applied recursively).
- An instance of a SQL collection type is rendered as a JSON array whose element values are taken from the collection element values (to which JSON generation is applied recursively).
- A `VARCHAR2`, `CLOB`, or `NVARCHAR` value is wrapped in double quotation marks (`"`), and characters are escaped when necessary to conform to the JSON standard.
- A numeric value (`NUMBER`, `BINARY_DOUBLE`, or `BINARY_FLOAT`) is rendered as a JSON number. (It is not quoted.)

The numeric values of positive and negative infinity, and values that are the undefined result of a numeric operation (`"not a number"` or `NaN`), cannot be expressed as JSON numbers. They are instead rendered as JSON strings: `"Inf"`, `"-Inf"`, and `"NaN"`, respectively.

- A `RAW` or `BLOB` value is rendered as a hexadecimal JSON string (with double quotation marks, `"`).
- A time-related value (`DATE`, `TIMESTAMP`, `TIMESTAMP WITH TIME ZONE`, `TIMESTAMP WITH LOCAL TIME ZONE`, `INTERVAL YEAR TO MONTH`, or `INTERVAL DAY TO SECOND`) is rendered in a supported ISO 8601 format, and the result is enclosed in double quotation marks (`"`).
- A `BOOLEAN` PL/SQL value is rendered as JSON `true` or `false`. (It is not quoted.)
- A `NULL` value is rendered as JSON `null`, regardless of the `NULL` data type. (It is not quoted.)

 **Note:**

For data types `CLOB` and `BLOB`, an empty instance is distinguished from `NULL` and is rendered as an empty JSON string (`""`). But for data types `VARCHAR2`, `NVARCHAR2`, and `RAW`, Oracle SQL treats an empty (zero-length) value as `NULL`, so do *not* expect such a value to be rendered as a JSON string.

The format of an input argument can affect the format of the data that is returned by the function. In particular, if an input is declared or automatically determined to be of format `JSON` then it is treated as JSON data when computing the return value. [Example 21-1](#) illustrates this — it explicitly uses `FORMAT JSON` to interpret the SQL string `"true"` as the JSON Boolean value `true`.

Example 21-1 Declaring an Input Value To Be JSON

This example specifies `FORMAT JSON` for SQL string values `'true'` and `'false'`, in order that the JSON Boolean values `true` and `false` are used. Without specifying `FORMAT JSON`, the values of field `hasCommission` would be the JSON *string* values `"true"` and `"false"`, not the JSON *Boolean* values `true` and `false`.

```
SELECT json_object('name'           VALUE first_name || ' ' || last_name,
                  'hasCommission' VALUE
                  CASE WHEN commission_pct IS NULL THEN 'false' ELSE
'true'
                  END FORMAT JSON)
FROM employees WHERE first_name LIKE 'W%';

JSON_OBJECT('NAME' ISFIRST_NAME || ' ' || LAST_NAME, '
-----
{"name": "William Gietz", "hasCommission": false}
{"name": "William Smith", "hasCommission": true}
{"name": "Winston Taylor", "hasCommission": false}
```

Related Topics

- [Overview of SQL/JSON Generation Functions](#)
You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.
- [SQL/JSON Function JSON_OBJECT](#)
SQL/JSON function `json_object` constructs JSON objects from the results of evaluating its argument SQL expressions.
- [SQL/JSON Function JSON_ARRAY](#)
SQL/JSON function `json_array` constructs a JSON array from the results of evaluating its argument SQL expressions.

 **See Also:**

olink:SQLRF-GUID-7B72E154-677A-4342-A1EA-C74C1EA928E6#GUID-7B72E154-677A-4342-A1EA-C74C1EA928E6

21.3 SQL/JSON Function JSON_OBJECT

SQL/JSON function `json_object` constructs JSON objects from the results of evaluating its argument SQL expressions.

It can accept any number of arguments, each of which is one of the following:

- An explicit field name–value pair. Example: `answer : 42`.

A *name–value* pair argument specifies an object member for the generated JSON object (except when the value expression evaluates to SQL NULL and the ABSENT ON NULL clause applies). The name and value are SQL expressions. The *name* expression must evaluate to a SQL *string*. The *value* expression must evaluate to a SQL value that can be rendered as a JSON value. The name and value expressions are separated by keyword **VALUE** or a colon (:).

 **Note:**

Some client drivers might try to scan query text and identify bind variables before sending the query to the database. In some such cases a colon as name–value separator in `json_object` might be misinterpreted as introducing a bind variable. You can use keyword **VALUE** as the separator to avoid this problem (`'Name' VALUE Diderot`), or you can simply enclose the value part of the pair in parentheses: `'Name' : (Diderot)`.

- A relational column name, possibly preceded by a table name or alias, or a view name followed by a dot (`.`). Example: `t1.address`.

In this case, for a given row of data, the JSON-object member specified by the column-name argument has the column name as its field name and the column value as the field value.

Regardless of whether it is quoted, the column name you provide is interpreted *case-sensitively*. For example, if you use `Email` as a column-name argument then the data in column `EMAIL` is used to produce object members with field name `Email` (not `EMAIL`).

- A table name or alias, or a view name, followed by a dot and an asterisk *wildcard* (`.*`). Example: `t1.*`. (The name or alias can also be prefixed by a database schema name, as in `myschema.t1.*`.)

In this case, all columns of the table or view are used as input. Each is handled as if it were named explicitly. In particular, the column names are interpreted *case-sensitively*.

Alternatively, `json_object` accepts a *single* argument that is one of the following:

- An instance of a user-defined SQL object-type. Example:
`json_object(my_sql_object_42).`

In this case, the resulting JSON-object field names are taken from the SQL object attribute names, and their values are taken from the SQL object attribute values (to which JSON generation is applied recursively).

- An asterisk *wildcard* (*). Example: `json_object(*)`.

The wildcard acts as a shortcut to explicitly specifying *all* of the columns of a table or view, to produce the object members. The resulting JSON-object field names are the *uppercase* column names. You can use a wildcard with a table, a view, or a table alias, which is understood from the FROM list. The columns can be of any SQL data type.

Note the difference between this case (`json_object(*)`) and the case described above, where the asterisk is preceded by an explicit table or view name (or table alias), followed by a dot: `json_object(t.*)`. In the `json_object(*)` case, the column names are *not* interpreted case-sensitively.

Another way of describing the use of asterisk wildcards with `json_object` is to say that it follows what is allowed for wildcards in a SQL SELECT list.

Just as for SQL/JSON condition `is json`, you can use keywords `STRICT` and `WITH UNIQUE KEYS` with functions `json_object` and `json_objectagg`. The behavior for each is the same as for `is json`.

Example 21-2 Using Name–Value Pairs with JSON_OBJECT

This example constructs a JSON object for each employee of table `hr.employees` (from standard database schema `HR`) whose salary is greater than 15000.

It passes explicit name–value pairs to specify the members of the JSON object. The object includes, as the value of its field `contactInfo`, an object with fields `mail` and `phone`.

Because the return value of `json_object` is JSON data, `FORMAT JSON` is deduced for the input format of field `contactInfo` — the explicit `FORMAT JSON` here is not needed.

```
SELECT json_object('id'           : employee_id,
                  'name'         : first_name || ' ' || last_name,
                  'hireDate'     : hire_date,
                  'pay'          : salary,
                  'contactInfo'  : json_object('mail' : email,
                                             'phone' : phone_number
                                             FORMAT JSON))
FROM hr.employees
WHERE salary > 15000;
```

-- The query returns rows such as this (pretty-printed here for clarity):

```
{ "id"           : 101,
  "name"         : "Neena Kochhar",
  "hireDate"     : "21-SEP-05",
  "pay"          : 17000,
  "contactInfo" : { "mail" : "NKOCHHAR",
                    "phone" : "515.123.4568" } }
```

Example 21-3 Using Column Names with JSON_OBJECT

This example constructs a JSON object for the employee whose `employee_id` is 101. The fields produced are named after the columns, but case-sensitively.

```
SELECT json_object(last_name,
                  hire_date,
                  salary,
                  'contactInfo' : json_object(email, phone_number)
                  FORMAT JSON)
FROM hr.employees
WHERE employee_id = 101;
```

-- The query returns rows such as this (pretty-printed here for clarity):

```
{ "last_name"   : "Neena Kochhar",
  "hire-date"  : "21-SEP-05",
  "salary"     : 17000,
  "contactInfo" : { "email"       : "NKOCHHAR",
                    "phone_number" : "515.123.4568" } }
```

Example 21-4 Using a Wildcard (*) with JSON_OBJECT

This example constructs a JSON object for each employee whose salary is greater than 15000. Each column of table `employees` is used to construct one object member, whose field name is the (uppercase) column name. Note that a SQL `NULL` value results in a JSON field value of `null`.

```
SELECT json_object(*)
FROM hr.employees
WHERE salary > 15000;
```

-- The query returns rows such as this (pretty-printed here for clarity):

```
JSON_OBJECT(*)
-----
{"EMPLOYEE_ID":100,
 "FIRST_NAME":"Steven",
 "LAST_NAME":"King",
 "EMAIL":"SKING",
 "PHONE_NUMBER":"515.123.4567",
 "HIRE_DATE":"2003-06-17T00:00:00",
 "JOB_ID":"AD_PRES",
 "SALARY":24000,
 "COMMISSION_PCT":null,
 "MANAGER_ID":null,
 "DEPARTMENT_ID":90}

{"EMPLOYEE_ID":101,
 "FIRST_NAME":"Neena",
 "LAST_NAME":"Kochhar",
 "EMAIL":"NKOCHHAR",
 "PHONE_NUMBER":"515.123.4568",
```

```

"HIRE_DATE": "2005-09-21T00:00:00",
"JOB_ID": "AD_VP",
"SALARY": 17000,
"COMMISSION_PCT": null,
"MANAGER_ID": 100,
"DEPARTMENT_ID": 90}

{"EMPLOYEE_ID": 102,
 "FIRST_NAME": "Lex",
 "LAST_NAME": "De Haan",
 "EMAIL": "LDEHAAN",
 "PHONE_NUMBER": "515.123.4569",
 "HIRE_DATE": "2001-01-13T00:00:00",
 "JOB_ID": "AD_VP",
 "SALARY": 17000,
 "COMMISSION_PCT": null,
 "MANAGER_ID": 100,
 "DEPARTMENT_ID": 90}

```

Example 21-5 Using JSON_OBJECT With ABSENT ON NULL

This example queries table `hr.locations` from standard database schema `HR` to create JSON objects with fields `city` and `province`.

The default `NULL`-handling behavior for `json_object` is `NULL ON NULL`.

In order to prevent the creation of a field with a null JSON value, this example uses `ABSENT ON NULL`. The `NULL SQL` value for column `state_province` when column `city` has value `'Singapore'` means that no `province` field is created for that location.

```

SELECT JSON_OBJECT('city'      VALUE city,
                  'province' : state_province ABSENT ON NULL)
FROM hr.locations
WHERE city LIKE 'S%';

JSON_OBJECT('CITY' IS CITY, 'PROVINCE' IS STATE_PROVINCE ABSENT ON NULL)
-----
{"city": "Southlake", "province": "Texas"}
{"city": "South San Francisco", "province": "California"}
{"city": "South Brunswick", "province": "New Jersey"}
{"city": "Seattle", "province": "Washington"}
{"city": "Sydney", "province": "New South Wales"}
{"city": "Singapore"}
{"city": "Stretford", "province": "Manchester"}
{"city": "Sao Paulo", "province": "Sao Paulo"}

```

Example 21-6 Using a User-Defined Object-Type Instance with JSON_OBJECT

This example creates table `po_ship` with column `shipping` of object type `shipping_t`. (It uses SQL/JSON function `json_value` to construct the `shipping_t` instances from JSON data — see [Example 16-3](#).)

It then uses `json_object` to generate JSON objects from the SQL object-type instances in column `po_ship.shipping`.

(The query output is shown pretty-printed here, for clarity.)

```
CREATE TABLE po_ship
AS SELECT json_value(po_document, '$.ShippingInstructions'
                    RETURNING shipping_t)
       shipping
FROM j_purchaseorder;

DESCRIBE po_ship;

Name          Null?    Type
-----
SHIPPING      SHIPPING_T

SELECT json_object(shipping) from po_ship;

JSON_OBJECT(SHIPPING)
-----
{"NAME":"Alexis Bull",
 "ADDRESS":{"STREET":"200 Sporting Green",
            "CITY":"South San Francisco"}}
{"NAME":"Sarah Bell",
 "ADDRESS":{"STREET":"200 Sporting Green",
            "CITY":"South San Francisco"}}
```

Related Topics

- [Overview of SQL/JSON Generation Functions](#)
You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.
- [Handling of Input Values For SQL/JSON Generation Functions](#)
The SQL/JSON generation functions take SQL values as input and render them as JSON values inside the SQL value that is returned. How the input values are rendered as JSON depends on their SQL data type.

See Also:

- *Oracle Database SQL Language Reference* for information about the `select_list` syntax
- *Oracle Database SQL Language Reference* for information about SQL/JSON function `json_object`
- *Oracle Database SQL Language Reference* for SQL identifier syntax

21.4 SQL/JSON Function JSON_ARRAY

SQL/JSON function `json_array` constructs a JSON array from the results of evaluating its argument SQL expressions.

In the simplest case, the evaluated arguments you provide to `json_array` are SQL values that are converted to JSON values as the JSON array elements. The resulting array has an element for each argument you provide (except when an argument expression evaluates to SQL NULL and the `ABSENT ON NULL` clause applies). Array element order is the same as the argument order.

There are several kinds of SQL values that you can use as an argument to `json_array`, including scalar, collection instance, and user-defined object-type instance.

Example 21-7 Using JSON_ARRAY to Construct a JSON Array

This example constructs a JSON object for each employee job in database table `hr.jobs` (from standard database schema `HR`). The fields of the objects are the job title and salary range. The salary range (field `salaryRange`) is an array of two numeric values, the minimum and maximum salaries for the job. These values are taken from SQL columns `min_salary` and `max_salary`.

```
SELECT json_object('title'          VALUE job_title,
                  'salaryRange' VALUE json_array(min_salary,
max_salary))
FROM jobs;

JSON_OBJECT('TITLE' ISJOB_TITLE, 'SALARYRANGE' ISJSON_ARRAY(MIN_SALARY,MAX_
SALARY))
```

```
-----
-----
{"title":"President","salaryRange":[20080,40000]}
{"title":"Administration Vice President","salaryRange":[15000,30000]}
{"title":"Administration Assistant","salaryRange":[3000,6000]}
{"title":"Finance Manager","salaryRange":[8200,16000]}
{"title":"Accountant","salaryRange":[4200,9000]}
{"title":"Accounting Manager","salaryRange":[8200,16000]}
{"title":"Public Accountant","salaryRange":[4200,9000]}
{"title":"Sales Manager","salaryRange":[10000,20080]}
{"title":"Sales Representative","salaryRange":[6000,12008]}
{"title":"Purchasing Manager","salaryRange":[8000,15000]}
{"title":"Purchasing Clerk","salaryRange":[2500,5500]}
{"title":"Stock Manager","salaryRange":[5500,8500]}
{"title":"Stock Clerk","salaryRange":[2008,5000]}
{"title":"Shipping Clerk","salaryRange":[2500,5500]}
{"title":"Programmer","salaryRange":[4000,10000]}
{"title":"Marketing Manager","salaryRange":[9000,15000]}
{"title":"Marketing Representative","salaryRange":[4000,9000]}
{"title":"Human Resources Representative","salaryRange":[4000,9000]}
{"title":"Public Relations Representative","salaryRange":[4500,10500]}
```

Related Topics

- [Overview of SQL/JSON Generation Functions](#)

You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.

- [Handling of Input Values For SQL/JSON Generation Functions](#)
The SQL/JSON generation functions take SQL values as input and render them as JSON values inside the SQL value that is returned. How the input values are rendered as JSON depends on their SQL data type.
- [SQL/JSON Function JSON_OBJECT](#)
SQL/JSON function `json_object` constructs JSON objects from the results of evaluating its argument SQL expressions.

**See Also:**

Oracle Database SQL Language Reference for information about SQL/JSON function `json_array`

21.5 SQL/JSON Function JSON_OBJECTAGG

SQL/JSON function `json_objectagg` constructs a JSON object by aggregating information from multiple rows of a grouped SQL query as the object members.

Unlike the case for SQL/JSON function `json_object`, where the number of members in the resulting object directly reflects the number of arguments, for `json_objectagg` the size of the resulting object reflects the current queried data. It can thus vary, depending on the data that is queried.

Example 21-8 Using JSON_OBJECTAGG to Construct a JSON Object

This example constructs a single JSON object from table `hr.departments` (from standard database schema `HR`) using field names taken from column `department_name` and field values taken from column `department_id`.

```
SELECT json_objectagg(department_name VALUE department_id) FROM
departments;
```

```
-- The returned object is pretty-printed here for clarity.
-- The order of the object members is arbitrary.
```

```
JSON_OBJECTAGG(DEPARTMENT_NAME|SDEPARTMENT_ID)
```

```
-----
{"Administration":      10,
 "Marketing":           20,
 "Purchasing":          30,
 "Human Resources":     40,
 "Shipping":            50,
 "IT":                  60,
 "Public Relations":    70,
 "Sales":               80,
 "Executive":           90,
 "Finance":             100,
 "Accounting":          110,
 "Treasury":            120,
 "Corporate Tax":       130,
 "Control And Credit":  140,
 "Shareholder Services": 150,
```

```

"Benefits":          160,
"Manufacturing":    170,
"Construction":     180,
"Contracting":      190,
"Operations":       200,
"IT Support":       210,
"NOC":              220,
"IT Helpdesk":      230,
"Government Sales": 240,
"Retail Sales":     250,
"Recruiting":       260,
"Payroll":          270}

```

Related Topics

- [Overview of SQL/JSON Generation Functions](#)

You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.

See Also:

Oracle Database SQL Language Reference for information about SQL/JSON function `json_objectagg`

21.6 SQL/JSON Function JSON_ARRAYAGG

SQL/JSON function `json_arrayagg` constructs a JSON array by aggregating information from multiple rows of a grouped SQL query as the array elements. The order of array elements reflects the query result order, by default, but you can use the `ORDER BY` clause to impose array element order.

Unlike the case for SQL/JSON function `json_array`, where the number of elements in the resulting array directly reflects the number of arguments, for `json_arrayagg` the size of the resulting array reflects the current queried data. It can thus vary, depending on the data that is queried.

Example 21-9 Using JSON_ARRAYAGG to Construct a JSON Array

This example constructs a JSON object for each employee of table `hr.employees` (from standard database schema `HR`) who is a manager in charge of at least six employees. The objects have fields for the manager id number, manager name, number of employees reporting to the manager, and id numbers of those employees.

The order of the employee id numbers in the array is determined by the `ORDER BY` clause for `json_arrayagg`. The default direction for `ORDER BY` is `ASC` (ascending). The array elements, which are numeric, are in ascending numerical order.

```

SELECT json_object('id'          VALUE mgr.employee_id,
                  'manager'     VALUE (mgr.first_name || ' ' ||
mgr.last_name),
                  'numReports'  VALUE count(rpt.employee_id),

```



```

        'reports'    VALUE json_arrayagg(rpt.employee_id
                                      ORDER BY
rpt.employee_id))
FROM employees mgr, employees rpt
WHERE mgr.employee_id = rpt.manager_id
GROUP BY mgr.employee_id, mgr.last_name, mgr.first_name
HAVING count(rpt.employee_id) > 6;

```

-- The returned object is pretty-printed here for clarity.

```

JSON_OBJECT('ID' ISMGR.EMPLOYEE_ID, 'MANAGER' VALUE(MGR.FIRST_NAME || ' ' ||
MGR.LAST_NAME)

```

```

{"id":      100,
 "manager": "Steven King",
 "numReports": 14,
 "reports": [101,102,114,120,121,122,123,124,145,146,147,148,149,201]}

```

```

{"id":      120,
 "manager": "Matthew Weiss",
 "numReports": 8,
 "reports": [125,126,127,128,180,181,182,183]}

```

```

{"id":      121,
 "manager": "Adam Fripp",
 "numReports": 8,
 "reports": [129,130,131,132,184,185,186,187]}

```

```

{"id":      122,
 "manager": "Payam Kaufling",
 "numReports": 8,
 "reports": [133,134,135,136,188,189,190,191]}

```

```

{"id":      123,
 "manager": "Shanta Vollman",
 "numReports": 8,
 "reports": [137,138,139,140,192,193,194,195]}

```

```

{"id":      124,
 "manager": "Kevin Mourgos",
 "numReports": 8,
 "reports": [141,142,143,144,196,197,198,199]}

```

Related Topics

- [Overview of SQL/JSON Generation Functions](#)
You can use SQL/JSON functions `json_object`, `json_array`, `json_objectagg`, and `json_arrayagg` to construct JSON data from non-JSON data in the database. The JSON data is returned as a SQL value.

 **See Also:**

Oracle Database SQL Language Reference for information about SQL/JSON function `json_arrayagg`

Part VI

PL/SQL Object Types for JSON

You can use PL/SQL object types for JSON to read and write multiple fields of a JSON document. This can increase performance, in particular by avoiding multiple parses and serializations of the data.

- [Overview of PL/SQL Object Types for JSON](#)
PL/SQL object types allow fine-grained programmatic construction and manipulation of In-Memory JSON data. You can introspect it, modify it, and serialize it back to textual JSON data.
- [Using PL/SQL Object Types for JSON](#)
Some examples of using PL/SQL object types for JSON are presented.

Overview of PL/SQL Object Types for JSON

PL/SQL object types allow fine-grained programmatic construction and manipulation of In-Memory JSON data. You can introspect it, modify it, and serialize it back to textual JSON data.

The principal PL/SQL JSON object types are `JSON_ELEMENT_T`, `JSON_OBJECT_T`, `JSON_ARRAY_T`, and `JSON_SCALAR_T`. Another, less used object type is `JSON_KEY_LIST`, which is a varray of `VARCHAR2(4000)`. Object types are also called abstract data types (ADTs).

These JSON object types provide an In-Memory, hierarchical (tree-like), programmatic representation of JSON data that is stored in the database.¹

You can use the object types to programmatically manipulate JSON data in memory, to do things such as the following:

- Check the structure, types, or values of existing JSON data. For example, check whether the value of a given object field satisfies certain conditions.
- Transform existing JSON data. For example, convert address or phone-number formats to follow a particular convention.
- Create JSON data using programming rules that match the characteristics of whatever the data represents. For example, if a product to be represented as a JSON object is flammable then include fields that represent safety information.

PL/SQL object-type instances are *transient*. To store the information they contain persistently, you must serialize them to `VARCHAR2` or LOB data, which you can then store in a database table or marshal to a database client such as Java Database Connectivity (JDBC).

You *construct* an object-type instance in memory either all at once, by parsing JSON text, or piecemeal, starting with an empty object or array instance and adding object members or array elements to it.

An unused object-type instance is automatically garbage-collected; you cannot, and need not, free up the memory used by an instance that you no longer need.

Typically, after you have constructed a PL/SQL object-type instance and perhaps made use of it programmatically in various ways, you *serialize* it to an instance of data type `VARCHAR2`, `CLOB`, or `BLOB`. That is, you convert the transient representation of JSON data in memory to a persistent representation in the database. (Alternatively, you might serialize it only as text to be printed out.)

¹ This is similar to what is available for XML data using the Document Object Model (DOM), a language-neutral and platform-neutral object model and API for accessing the structure of XML documents that is recommended by the World Wide Web Consortium (W3C).

Relations Among the JSON Object Types

Type `JSON_ELEMENT_T` is the supertype of the other JSON object types: each of them extends it as a subtype. Subtypes `JSON_OBJECT_T` and `JSON_ARRAY_T` are used for JSON objects and arrays, respectively. Subtype `JSON_SCALAR_T` is used for scalar JSON values: strings, numbers, the Boolean values `true` and `false`, and the value `null`.

You can construct an instance of type `JSON_ELEMENT_T` only by parsing JSON text. Parsing creates a `JSON_ELEMENT_T` instance, which is an In-Memory representation of the JSON data. You cannot construct an empty instance of type `JSON_ELEMENT_T` or type `JSON_SCALAR_T`.

Types `JSON_OBJECT_T` and `JSON_ARRAY_T` each have a constructor function of the same name as the type, which you can use to construct an instance of the type: an empty (In-Memory) representation of a JSON object or array, respectively. You can then fill this object or array as needed, adding object members or array elements, represented by PL/SQL object-type instances.

You can cast an instance of `JSON_ELEMENT_T` to a subtype instance, using PL/SQL function `treat`. For example, `treat(elt AS JSON_OBJECT_T)` casts instance `elt` as a JSON object (instance of `JSON_OBJECT_T`).

Parsing Function

Static function `parse` accepts an instance of type `VARCHAR2`, `CLOB`, or `BLOB` as argument, which it parses as JSON text to return an instance of type `JSON_ELEMENT_T`, `JSON_OBJECT_T`, or `JSON_ARRAY_T`.

Serialization Methods

Parsing accepts input JSON data as text and returns an instance of a PL/SQL JSON object type. Serialization does essentially the opposite: you apply it to a PL/SQL object representation of JSON data and it returns a textual representation of that object. The serialization methods have names that start with prefix `to_`. For example, method `to_string()` returns a string (`VARCHAR2`) representation of the JSON object-type instance you apply it to.

Most serialization methods are member functions. For serialization as a `CLOB` or `BLOB` instance, however, there are two forms of the methods: a member *function* and a *member procedure*. The member function accepts no arguments. It creates a temporary LOB as the serialization destination. The member procedure accepts a LOB IN OUT argument (`CLOB` instance for method `to_clob`, `BLOB` for method `to_blob`). You can thus pass it the LOB (possibly empty) that you want to use for the serialized representation.

Getter and Setter Methods

Types `JSON_OBJECT_T` and `JSON_ARRAY_T` have getter and setter methods, which obtain and update, respectively, the values of a given object field or a given array element position.

There are two kinds of *getter* method:

- Method `get()` returns a reference to the original object to which you apply it, as an instance of type `JSON_ELEMENT_T`. That is, the object to which you apply it is

passed by reference: If you then modify the returned `JSON_ELEMENT_T` instance, your modifications apply to the original object to which you applied `get()`.

- Getter methods whose names have the prefix `get_` return a *copy* of any data that is targeted within the object or array to which they are applied. That data is *passed by value*, not reference.

For example, if you apply method `get_string()` to a `JSON_OBJECT_T` instance, passing a given field as argument, it returns a copy of the string that is the value of that field. If you apply `get_string()` to a `JSON_ARRAY_T` instance, passing a given element position as argument, it returns a copy of the string at that position in the array.

Like the serialization methods, most getter methods are member functions. But methods `get_clob()` and `get_blob()`, which return the value of a given object field or the element at a given array position as a `CLOB` or `BLOB` instance, have two forms (like the serialization methods `to_clob()` and `to_blob()`): a member *function* and a member *procedure*. The member function accepts no argument other than the targeted object field or array position. It creates and returns a temporary LOB instance. The member procedure accepts also a `LOB IN OUT` argument (`CLOB` for `get_clob`, `BLOB` for `get_blob`). You can thus pass it the (possibly empty) LOB instance to use.

The *setter* methods are `put()`, `put_null()`, and (for `JSON_ARRAY_T` only) `append()`. These update the object or array instance to which they are applied, setting the value of the targeted object field or array element. Note: The setter methods *modify the existing instance*, instead of returning a modified copy of it.

Method `append()` adds a new element at the end of the array instance. Method `put_null()` sets an object field or array element value to `JSON null`.

Method `put()` requires a second argument (besides the object field name or array element position), which is the new value to set. For an array, `put()` also accepts an optional third argument, `OVERWRITE`. This is a `BOOLEAN` value (default `FALSE`) that says whether to *replace an existing value* at the given position.

- If the object already has a field of the same name then `put()` *replaces that value* with the new value.
- If the array already has an element at the given position then, by default, `put()` shifts that element and any successive elements forward (incrementing their positions by one) to make room for the new element, which is placed at the given position. But if optional argument `OVERWRITE` is present and is `TRUE`, then the existing element at the given position is simply *replaced* by the new element.

Introspection Methods

Type `JSON_ELEMENT_T` has introspection methods that you can use to determine whether an instance is a JSON object, array, scalar, string, number, or Boolean, or whether it is the JSON value `true`, `false`, or `null`. The names of these methods begin with prefix `is_`. They are predicates, returning a `BOOLEAN` value.

It also has introspection method `get_size()`, which returns the number of members of a `JSON_OBJECT_T` instance and the number of elements of a `JSON_ARRAY_T` instance (it returns 1 for a `JSON_SCALAR_T` instance).

Type `JSON_ELEMENT_T` also has introspection methods `is_date()` and `is_timestamp()`, which test whether an instance represents a date or timestamp. JSON has no native types for dates or timestamps; these are typically representing using JSON strings.

But if a `JSON_ELEMENT_T` instance is constructed using SQL data of SQL data type `DATE` or `TIMESTAMP` then this type information is kept for the PL/SQL object representation.

Date and timestamp data is represented using PL/SQL object type `JSON_SCALAR_T`, whose instances you cannot construct directly. You can, however, add such a value to an object (as a field value) or an array (as an element) using method `put()`. Retrieving it using method `get()` returns a `JSON_SCALAR_T` instance.

Types `JSON_OBJECT_T` and `JSON_ARRAY_T` have introspection method `get_type()`, which returns the JSON type of the targeted object field or array element (as a `VARCHAR2` instance). Type `JSON_OBJECT_T` also has introspection methods `has()`, which returns `TRUE` if the object has a field of the given name, and `get_keys()`, which returns an instance of PL/SQL object type `JSON_KEY_LIST`, which is a varray of type `VARCHAR2(4000)`. The varray contains the names of the fields ² present in the given `JSON_OBJECT_T` instance.

Other Methods

Types `JSON_OBJECT_T` and `JSON_ARRAY_T` have the following methods:

- `remove()` — Remove the object member with the given field or the array element at the given position.
- `clone()` — Create and return a (deep) copy of the object or array to which the method is applied. Modifying any part of this copy has no effect on the original object or array.

Type `JSON_OBJECT_T` has method `rename_key()`, which renames a given object field.² If the new name provided already names an existing field then an error is raised.

Related Topics

- [Using PL/SQL Object Types for JSON](#)
Some examples of using PL/SQL object types for JSON are presented.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_ARRAY_T`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_ELEMENT_T`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_OBJECT_T` and `JSON_KEY_LIST`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_SCALAR_T`

² An object field is sometimes called an object “key”.

23

Using PL/SQL Object Types for JSON

Some examples of using PL/SQL object types for JSON are presented.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_ARRAY_T`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_ELEMENT_T`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_OBJECT_T`
- *Oracle Database PL/SQL Packages and Types Reference* for information about `JSON_KEY_LIST`

Example 23-1 Constructing and Serializing an In-Memory JSON Object

This example uses function `parse` to parse a string of JSON data that represents a JSON object with one field, `name`, creating an instance `je` of object type `JSON_ELEMENT_T`. This instance is tested to see if it represents an object, using introspection method (predicate) `is_object()`.

If it represents an object (the predicate returns `TRUE` for `je`), it is cast to an instance of `JSON_OBJECT_T` and assigned to variable `jo`. Method `put()` for object type `JSON_OBJECT_T` is then used to add object field `price` with value `149.99`.

Finally, `JSON_ELEMENT_T` instance `je` (which is the same data in memory as `JSON_OBJECT_T` instance `jo`) is serialized to a string using method `to_string()`, and this string is printed out using procedure `DBMS_OUTPUT.put_line`. The result printed out shows the updated object as `{"name": "Radio-controlled plane", "price": 149.99}`.

The updated transient object `je` is serialized here only to be printed out; the resulting text is not stored in the database. Sometime after the example code is executed, the memory allocated for object-type instances `je` and `jo` is reclaimed by the garbage collector.

```
DECLARE
  je JSON_ELEMENT_T;
  jo JSON_OBJECT_T;
BEGIN
  je := JSON_ELEMENT_T.parse('{"name": "Radio controlled plane"}');
  IF (je.is_Object) THEN
    jo := treat(je AS JSON_OBJECT_T);
    jo.put('price', 149.99);
  END IF;
  DBMS_OUTPUT.put_line(je.to_string);
```



```
END;
/
```

Example 23-2 Using Method GET_KEYS() to Obtain a List of Object Fields

PL/SQL method `get_keys()` is defined for PL/SQL object type `JSON_OBJECT_T`. It returns an instance of PL/SQL object type `JSON_KEY_LIST`, which is a varray of `VARCHAR2(4000)`. The varray contains all of the field names for the given `JSON_OBJECT_T` instance.

This example iterates through the fields returned by `get_keys()`, adding them to an instance of PL/SQL object type `JSON_ARRAY_T`. It then uses method `to_string()` to serialize that JSON array and then prints the resulting string.

```
DECLARE
  jo          JSON_OBJECT_T;
  ja          JSON_ARRAY_T;
  keys       JSON_KEY_LIST;
  keys_string VARCHAR2(100);
BEGIN
  ja := new JSON_ARRAY_T;
  jo := JSON_OBJECT_T.parse('{"name": "Beda",
                           "jobTitle": "codmonki",
                           "projects": ["json", "xml"]}');

  keys := jo.get_keys;
  FOR i IN 1..keys.COUNT LOOP
    ja.append(keys(i));
  END LOOP;
  keys_string := ja.to_string;
  DBMS_OUTPUT.put_line(keys_string);
END;
/
```

The printed output is `["name", "jobTitle", "projects"]`.

Example 23-3 Using Method PUT() to Update Parts of JSON Documents

This example updates each purchase-order document in JSON column `po_document` of table `j_purchaseorder`. It iterates over the JSON array `LineItems` in each document (variable `li_arr`), calculating the total price and quantity for each line-item object (variable `li_obj`), and it uses method `put()` to add these totals to `li_obj` as the values of new fields `totalQuantity` and `totalPrice`. This is done by user-defined function `add_totals`.

The `SELECT` statement here selects one of the documents that has been updated.

```
CREATE OR REPLACE FUNCTION add_totals(purchaseOrder IN VARCHAR2) RETURN
VARCHAR2 IS
  po_obj      JSON_OBJECT_T;
  li_arr      JSON_ARRAY_T;
  li_item     JSON_ELEMENT_T;
  li_obj      JSON_OBJECT_T;
  unitPrice   NUMBER;
  quantity    NUMBER;
  totalPrice  NUMBER := 0;
```

```

totalQuantity NUMBER := 0;
BEGIN
po_obj := JSON_OBJECT_T.parse(purchaseOrder);
li_arr := po_obj.get_Array('LineItems');
FOR i IN 0 .. li_arr.get_size - 1 LOOP
li_obj := JSON_OBJECT_T(li_arr.get(i));
quantity := li_obj.get_Number('Quantity');
unitPrice := li_obj.get_Object('Part').get_Number('UnitPrice');
totalPrice := totalPrice + (quantity * unitPrice);
totalQuantity := totalQuantity + quantity;
END LOOP;
po_obj.put('totalQuantity', totalQuantity);
po_obj.put('totalPrice', totalPrice);
RETURN po_obj.to_string;
END;
/

UPDATE j_purchaseorder SET (po_document) = add_totals(po_document);

SELECT po_document FROM j_purchaseorder po WHERE
po.po_document.PONumber = 1600;

```

That selects this updated document:

```

{"PONumber": 1600,
"Reference": "ABULL-20140421",
"Requestor": "Alexis Bull",
"User": "ABULL",
"CostCenter": "A50",
"ShippingInstructions": {"name": "Alexis Bull",
"Address": {"street": "200 Sporting Green",
"city": "South San Francisco",
"state": "CA",
"zipCode": 99236,
"country": "United States of
America"}},
"Phone": [{"type": "Office", "number":
"909-555-7307"},
{"type": "Mobile", "number":
"415-555-1234"}]},
"Special Instructions": null,
"AllowPartialShipment": true,
"LineItems": [{"ItemNumber": 1,
"Part": {"Description": "One Magic Christmas",
"UnitPrice": 19.95,
"UPCCode": 13131092899},
"Quantity": 9.0},
{"ItemNumber": 2,
"Part": {"Description": "Lethal Weapon",
"UnitPrice": 19.95,
"UPCCode": 85391628927},
"Quantity": 5.0}],
"totalQuantity": 14,
"totalPrice": 279.3}

```

Related Topics

- [Overview of PL/SQL Object Types for JSON](#)
PL/SQL object types allow fine-grained programmatic construction and manipulation of In-Memory JSON data. You can introspect it, modify it, and serialize it back to textual JSON data.
- [Updating a JSON Document with JSON Merge Patch](#)
You can use Oracle SQL function `json_mergepatch` to update specific portions of a JSON document. You pass it a JSON Merge Patch document, which specifies the changes to make to a specified JSON document. JSON Merge Patch is an IETF standard.

Part VII

GeoJSON Geographic Data

GeoJSON data is geographic JSON data. Oracle Spatial and Graph supports the use of GeoJSON objects to store, index, and manage GeoJSON data.

- [Using GeoJSON Geographic Data](#)
GeoJSON objects are JSON objects that represent geographic data. Examples are provided of creating GeoJSON data, indexing it, and querying it.

Using GeoJSON Geographic Data

GeoJSON objects are JSON objects that represent geographic data. Examples are provided of creating GeoJSON data, indexing it, and querying it.

GeoJSON Objects: Geometry, Feature, Feature Collection

GeoJSON uses JSON objects that represent various geometrical entities and combinations of these together with user-defined properties.

A **position** is an array of two or more spatial (numerical) coordinates, the first three of which generally represent longitude, latitude, and altitude.

A **geometry** object has a `type` field and (except for a geometry-collection object) a `coordinates` field, as shown in [Table 24-1](#).

A **geometry collection** is a geometry object with `type` `GeometryCollection`. Instead of a `coordinates` field it has a `geometries` field, whose value is an array of geometry objects other than `GeometryCollection` objects.

Table 24-1 GeoJSON Geometry Objects Other Than Geometry Collections

<code>type</code> Field	<code>coordinates</code> Field
<code>Point</code>	A position.
<code>MultiPoint</code>	An array of positions.
<code>LineString</code>	An array of two or more positions.
<code>MultiLineString</code>	An array of <code>LineString</code> arrays of positions.
<code>Polygon</code>	A <code>MultiLineString</code> , each of whose arrays is a <code>LineString</code> whose first and last positions coincide (are equivalent). If the array of a polygon contains more than one array then the first represents the outside polygon and the others represent holes inside it.
<code>MultiPolygon</code>	An array of <code>Polygon</code> arrays, that is, multidimensional array of positions.

A **feature** object has a `type` field of value `Feature`, a `geometry` field whose value is a geometric object, and a `properties` field whose value can be any JSON object.

A **feature collection** object has a `type` field of value `FeatureCollection`, and it has a `features` field whose value is an array of feature objects.

[Example 24-1](#) presents a feature-collection object whose `features` array has three features. The `geometry` of the first feature is of type `Point`; that of the second is of type `LineString`; and that of the third is of type `Polygon`.

Query and Index GeoJSON Data

You can use SQL/JSON query functions and conditions to examine GeoJSON data or to project parts of it as non-JSON data, including as Oracle Spatial and Graph

SDO_GEOMETRY object-type instances. This is illustrated in [Example 24-2](#), [Example 24-3](#), and [Example 24-5](#).

To improve query performance, you can create an Oracle Spatial and Graph index (type `MDSYS.SPATIAL_INDEX`) on function `json_value` applied to GeoJSON data. This is illustrated by [Example 24-4](#).

[Example 24-4](#) indexes only one particular element of an array of geometry features (the first element). A B-tree index on function `json_value` can target only a *scalar* value. To improve the performance of queries, such as that of [Example 24-3](#), that target any number of array elements, you can do the following:

- Create an on-statement, refreshable *materialized view* of the array data, and place that view *in memory*.
- Create a spatial index on the array data.

This is shown in [Example 24-6](#) and [Example 24-7](#).

SDO_GEOMETRY Object-Type Instances and Spatial Operations

You can convert Oracle Spatial and Graph `SDO_GEOMETRY` object-type instances to GeoJSON objects and GeoJSON objects to `SDO_GEOMETRY` instances.

You can use Oracle Spatial and Graph operations on `SDO_GEOMETRY` objects that you obtain from GeoJSON objects. For example, you can use operator `sdo_distance` in PL/SQL package `SDO_GEOM` to compute the minimum distance between two geometry objects. This is the distance between the closest two points or two segments, one point or segment from each object. This is illustrated by [Example 24-5](#).

JSON Data Guide Supports GeoJSON Data

A JSON data guide summarizes structural and type information contained in a set of JSON documents. If some of the documents contain GeoJSON data then that data is summarized in a data guide that you create using SQL aggregate function `json_dataguide`. If you use SQL function `json_dataguide` to create a view based on such a data guide, and you specify the formatting argument as `DBMS_JSON.GEOJSON` or `DBMS_JSON.GEOJSON+DBMS_JSON.PRETTY`, then a column that projects GeoJSON data from the document set is of SQL data type `SDO_GEOMETRY`.

See Also:

- *Oracle Spatial and Graph Developer's Guide* for information about using GeoJSON data with Oracle Spatial and Graph
- *Oracle Spatial and Graph Developer's Guide* for information about Oracle Spatial and Graph and `SDO_GEOMETRY` object type
- GeoJSON.org for information about GeoJSON
- *The GeoJSON Format Specification* for details about GeoJSON data

Example 24-1 A Table With GeoJSON Data

This example creates table `j_geo`, which has a column, `geo_doc` of GeoJSON documents.

Only one such document is inserted here. It contains a GeoJSON object of type `FeatureCollection`, and a `features` array of objects of type `Feature`. Those objects have geometry, respectively, of type `Point`, `LineString`, and `Polygon`.

```
CREATE TABLE j_geo
(id          VARCHAR2 (32) NOT NULL,
 geo_doc    VARCHAR2 (4000) CHECK (geo_doc IS JSON));

INSERT INTO j_geo
VALUES (1,
       '{"type"      : "FeatureCollection",
        "features"  : [{"type"      : "Feature",
                          "geometry" : {"type" : "Point",
                                         "coordinates" :
[-122.236111, 37.482778]}},
                          "properties" : {"Name" : Redwood City}}},
        {"type" : "Feature",
         "geometry" : {"type" : "LineString",
                       "coordinates" : [[102.0, 0.0],
                                         [103.0, 1.0],
                                         [104.0, 0.0],
                                         [105.0,
1.0]]},
         "properties" : {"prop0" : "value0",
                          "prop1" : 0.0}},
        {"type" : "Feature",
         "geometry" : {"type" : "Polygon",
                       "coordinates" : [[[100.0,
0.0],
                                         [101.0,
0.0],
                                         [101.0,
1.0],
                                         [100.0,
1.0],
                                         [100.0,
0.0]]]}},
        "properties" : {"prop0" : "value0",
                          "prop1" : {"this" :
"that"}}}]'});
```

Example 24-2 Selecting a geometry Object From a GeoJSON Feature As an SDO_GEOMETRY Instance

This example uses SQL/JSON function `json_value` to select the value of field `geometry` from the first element of array `features`. The value is returned as Oracle Spatial and Graph data, not as JSON data, that is, as an instance of PL/SQL object type `SDO_GEOMETRY`, not as a SQL string or LOB instance.

```
SELECT json_value(geo_doc, '$.features[0].geometry'
                RETURNING SDO_GEOMETRY
                ERROR ON ERROR)
FROM j_geo;
```

The value returned is this, which represents a point with longitude and latitude (coordinates) -122.236111 and 37.482778, respectively.

```
SDO_GEOMETRY(2001, 4326, SDO_POINT_TYPE(-122.236111, 37.482778, NULL),
NULL, NULL)
```

See Also:

Oracle Database SQL Language Reference for information about SQL/JSON function `json_value`

Example 24-3 Retrieving Multiple geometry Objects From a GeoJSON Feature As SDO_GEOMETRY

This example uses SQL/JSON function `json_table` to project the value of field `geometry` from *each* element of array `features`, as column `sdo_val` of a virtual table. The retrieved data is returned as `SDO_GEOMETRY`.

```
SELECT jt.*
FROM j_geo,
     json_table(geo_doc, '$.features[*]'
               COLUMNS (sdo_val SDO_GEOMETRY PATH '$.geometry')) jt;
```

See Also:

Oracle Database SQL Language Reference for information about SQL/JSON function `json_table`

The following three rows are returned for the query. The first represents the same Point as in [Example 24-2](#). The second represents the `LineString` array. The third represents the `Polygon`.

```
SDO_GEOMETRY(2001, 4326, SDO_POINT_TYPE(-122.236111, 37.482778, NULL),
NULL, NULL)
```

```
SDO_GEOMETRY(2002, 4326, NULL, SDO_ELEM_INFO_ARRAY(1, 2, 1),
SDO_ORDINATE_ARRAY(102, 0, 103, 1, 104, 0, 105, 1))
```

```
SDO_GEOMETRY(2003, 4326, NULL, SDO_ELEM_INFO_ARRAY(1, 1003, 1),
SDO_ORDINATE_ARRAY(100, 0, 101, 0, 101, 1, 100, 1, 100, 0))
```

The second and third elements of attribute `SDO_ELEM_INFO_ARRAY` specify how to interpret the coordinates provided by attribute `SDO_ORDINATE_ARRAY`. They show that the first row returned represents a *line string* (2) with straight segments (1), and the second row represents a *polygon* (2003) of straight segments (1).

Example 24-4 Creating a Spatial Index For Scalar GeoJSON Data

This example creates a `json_value` function-based index of type `MDSYS.SPATIAL_INDEX` on field `geometry` of the first element of array `features`. This can improve the performance of queries that use `json_value` to retrieve that value.

```
CREATE INDEX geo_first_feature_idx
  ON j_geo (json_value(geo_doc, '$.features[0].geometry'
                     RETURNING SDO_GEOMETRY))
  INDEXTYPE IS MDSYS.SPATIAL_INDEX;
```

Example 24-5 Using GeoJSON Geometry With Spatial Operators

This example selects the documents (there is only one in this table) for which the `geometry` field of the first `features` element is within 100 kilometers of a given point. The point is provided literally here (its `coordinates` are the longitude and latitude of San Francisco, California). The distance is computed from this point to each geometry object.

The query orders the selected documents by the calculated distance. The tolerance in meters for the distance calculation is provided in this query as the literal argument 100.

```
SELECT id,
       json_value(geo_doc, '$features[0].properties.Name') "Name",
       SDO_GEOM.sdo_distance(
         json_value(geo_doc, '$features[0].geometry') RETURNING
         SDO_GEOMETRY,
         SDO_GEOMETRY(2001,
                      4326,
                      SDO_POINT_TYPE(-122.416667, 37.783333, NULL),
                      NULL,
                      NULL),
         100, -- Tolerance in meters
         'unit=KM') "Distance in kilometers"
FROM   j_geo
WHERE  sdo_within_distance(
       json_value(geo_doc, '$.features[0].geometry' RETURNING
       SDO_GEOMETRY),
       SDO_GEOMETRY(2001,
                    4326,
                    SDO_POINT_TYPE(-122.416667, 37.783333, NULL),
                    NULL,
                    NULL),
       'distance=100 unit=KM')
= 'TRUE';
```

See Also:

Oracle Database SQL Language Reference for information about SQL/JSON function `json_value`

The query returns a single row:

ID	Name	Distance in kilometers
1	Redwood City	26.9443035

Example 24-6 Creating a Materialized View Over GeoJSON Data

```
CREATE OR REPLACE MATERIALIZED VIEW geo_doc_view
  BUILD IMMEDIATE
  REFRESH FAST ON STATEMENT WITH ROWID
  AS SELECT g.rowid, jt.*
     FROM j_geo g,
          json_table(geo_doc, '$.features[*]'
                    COLUMNS (sdo_val SDO_GEOMETRY PATH '$.geometry')) jt;
```

Example 24-7 Creating a Spatial Index on a Materialized View Over GeoJSON Data

This example first prepares for the creation of the spatial index by populating some spatial-indexing metadata. It then creates the index on the `SDO_GEOMETRY` column, `sdo_val`, of materialized view `geo_doc_view`, which is created in [Example 24-6](#). Except for the view and column names, the code for populating the indexing metadata is fixed — use it each time you need to create a spatial index on a materialized view.

```
-- Populate spatial-indexing metadata

INSERT INTO USER_SDO_GEOM_METADATA
  VALUES ('GEO_DOC_VIEW',
         'SDO_VAL',
         MDSYS.sdo_dim_array(
           MDSYS.sdo_dim_element('Longitude', -180, 180, 0.05),
           MDSYS.sdo_dim_element('Latitude', -90, 90, 0.05)),
         7
         4326);

-- Create spatial index on geometry column of materialized view

CREATE INDEX geo_all_features_idx ON geo_doc_view(sdo_val)
  INDEXTYPE IS MDSYS.SPATIAL_INDEX V2;
```

Related Topics

- [Creating a View Over JSON Data Using JSON_TABLE](#)
To improve query performance you can create a view over JSON data that you project to columns using SQL/JSON function `json_table`. To further improve query performance you can create a *materialized view* and place the JSON data *in memory*.
- [JSON Data-Guide Fields](#)
The predefined fields of a JSON data guide are described. They include JSON Schema fields (keywords) and Oracle-specific fields.

Part VIII

Performance Tuning for JSON

To tune query performance you can index JSON fields in several ways, store their values in the In-Memory Column Store (IM column store), or expose them as non-JSON data using materialized views.

- [Overview of Performance Tuning for JSON](#)
Which performance-tuning approaches you take depend on the needs of your application. Some use cases and recommended solutions are outlined here.
- [Indexes for JSON Data](#)
You can index JSON data as you would any data of the type you use to store it. In addition, you can define a JSON search index, which is useful for both ad hoc structural queries and full-text queries.
- [In-Memory JSON Data](#)
A column of JSON data can be stored in the In-Memory Column Store (IM column store) to improve query performance.

Overview of Performance Tuning for JSON

Which performance-tuning approaches you take depend on the needs of your application. Some use cases and recommended solutions are outlined here.

The use cases can be divided into two classes: searching for or accessing data based on values of JSON fields that occur (1) at most once in a given document or (2) possibly more than once.

Queries That Access the Values of Fields That Occur at Most Once in a Given Document

You can tune the performance of such queries in the same ways as for non-JSON data. The choices of which JSON fields to define virtual columns for or which to index, whether to place the table containing your JSON data in the In-Memory Column Store (IM column store), and whether to create materialized views that project some of its fields are analogous to the non-JSON case.

However, in the case of JSON data it is generally *more* important to apply at least one such performance tuning than it is in the case non-JSON data. Without any such performance aid, it is typically more expensive to access a JSON field than it is to access (non-JSON) column data, because a JSON document must be traversed to locate the data you seek.

Create virtual columns from JSON fields or index JSON fields:

- If your queries use simple and highly selective search criteria, for a *single JSON field*:
 - Define a virtual column on the field.
You can often improve performance further by placing the table in the IM column store or creating an index on the virtual column.
 - Create a function-based index on the field using SQL/JSON function `json_value`.
- If your queries involve *more than one field*:
 - Define a virtual column on each of the fields.
You can often improve performance further by placing the table in the IM column store or creating a composite index on the virtual columns.
 - Create a composite function-based index on the fields using multiple invocations of SQL/JSON function `json_value`, one for each field.

Queries That Access the Values of Fields That Can Occur More Than Once in a Given Document

In particular, this is the case when you access fields that are contained within an array.

There are three techniques you can use to tune the performance of such queries:

- Place the table that contains the JSON data in the IM column store.

- Use a JSON search index.

This indexes all of the fields in a JSON document along with their values, including fields that occur inside arrays. The index can optimize any path-based search, including those using path expressions that include filters and full-text operators. The index also supports range-based searches on numeric values.

- Use a *materialized view* of non-JSON columns that are projected from JSON field values using SQL/JSON function `json_table`.

You can generate a separate row from each member of a JSON array, using the `NESTED PATH` clause with `json_table`.

A materialized view is typically used for optimizing SQL-based reporting and analytics for JSON content.

Indexes for JSON Data

You can index JSON data as you would any data of the type you use to store it. In addition, you can define a JSON search index, which is useful for both ad hoc structural queries and full-text queries.

- [Overview of Indexing JSON Data](#)
There is no dedicated SQL data type for JSON data, so you can index it in the usual ways. In addition, you can index it in a general way, with a JSON search index, for ad hoc structural queries and full-text queries.
- [How To Tell Whether a Function-Based Index for JSON Data Is Picked Up](#)
To determine whether a given query picks up a given function-based index, look for the index name in the execution plan for the query.
- [Creating Bitmap Indexes for SQL/JSON Condition JSON_EXISTS](#)
You can create a bitmap index for the value returned by `json_exists`. This is the right kind of index to use for `json_exists`, because there are only two possible return values for a condition (true and false).
- [Creating JSON_VALUE Function-Based Indexes](#)
You can create a function-based index for SQL/JSON function `json_value`. You can use the standard syntax for this, explicitly specifying function `json_value`, or you can use the simple dot-notation syntax. Indexes created in either of these ways can be used with both dot-notation queries and `json_value` queries.
- [Using a JSON_VALUE Function-Based Index with JSON_TABLE Queries](#)
An index created using `json_value` with `ERROR ON ERROR` can be used for a query involving `json_table`, if the `WHERE` clause refers to a column projected by `json_table`, and the effective SQL/JSON path that targets that column matches the indexed path expression.
- [Using a JSON_VALUE Function-Based Index with JSON_EXISTS Queries](#)
An index created using SQL/JSON function `json_value` with `ERROR ON ERROR` can be used for a query involving SQL/JSON condition `json_exists`, provided the query path expression has a filter expression that contains only a *path-expression comparison* or multiple such comparisons separated by `&&`.
- [Data Type Considerations for JSON_VALUE Indexing and Querying](#)
By default, SQL/JSON function `json_value` returns a `VARCHAR2` value. When you create a function-based index using `json_value`, unless you use a `RETURNING` clause to specify a different return data type, the index is not picked up for a query that expects a non-`VARCHAR2` value.
- [Indexing Multiple JSON Fields Using a Composite B-Tree Index](#)
To index multiple fields of a JSON object, you first create virtual columns for them. Then you create a composite B-tree index on the virtual columns.
- [JSON Search Index: Ad Hoc Queries and Full-Text Search](#)
A JSON search index is a *general* index. It can improve the performance of both (1) ad hoc structural queries, that is, queries that you might not anticipate or use regularly, and (2) full-text search. It is an Oracle Text index that is designed specifically for use with JSON data.

26.1 Overview of Indexing JSON Data

There is no dedicated SQL data type for JSON data, so you can index it in the usual ways. In addition, you can index it in a general way, with a JSON search index, for ad hoc structural queries and full-text queries.

You can index JSON data as you would any data of the type that you use to store it. In particular, you can use a B-tree index or a bitmap index for SQL/JSON function `json_value`, and you can use a bitmap index for SQL/JSON conditions `is json`, `is not json`, and `json_exists`.

(More generally, a bitmap index can be appropriate wherever the number of possible values for the function is small. For example, you can use a bitmap index for function `json_value` if the value is expected to be Boolean or otherwise one of a small number of string values.)

As always, such function-based indexing is appropriate for queries that target particular functions, which in the context of SQL/JSON functions means particular SQL/JSON *path expressions*. It is not very helpful for queries that are ad hoc, that is, arbitrary. Define a function-based index if you know that you will often query a particular path expression.

If you query in an ad hoc manner then define a **JSON search index**. This is a general index, *not targeted* to any specific path expression. It is appropriate for *structural* queries, such as looking for a JSON field with a particular value, and for *full-text* queries using SQL/JSON condition `json_textcontains`, such as looking for a particular word among various string values.

You can of course define both function-based indexes and a JSON search index for the same JSON column.

A JSON search index is an Oracle Text (full-text) index designed specifically for use with JSON data.

Note:

Oracle recommends that you use AL32UTF8 as the database character set. Automatic character-set conversion can take place when creating or applying an index. Such conversion can be lossy, which can mean that some data that you might expect to be returned by a query is not returned. See [Character Sets and Character Encoding for JSON Data](#).

Related Topics

- [Using GeoJSON Geographic Data](#)
GeoJSON objects are JSON objects that represent geographic data. Examples are provided of creating GeoJSON data, indexing it, and querying it.
- [JSON Search Index: Ad Hoc Queries and Full-Text Search](#)
A JSON search index is a *general* index. It can improve the performance of both (1) ad hoc structural queries, that is, queries that you might not anticipate or use regularly, and (2) full-text search. It is an Oracle Text index that is designed specifically for use with JSON data.

26.2 How To Tell Whether a Function-Based Index for JSON Data Is Picked Up

To determine whether a given query picks up a given function-based index, look for the index name in the execution plan for the query.

For example, given the index defined in [Example 26-4](#), an execution plan for the `json_value` query of [Example 12-1](#) references an index scan with index `po_num_id1`.

26.3 Creating Bitmap Indexes for SQL/JSON Condition `JSON_EXISTS`

You can create a bitmap index for the value returned by `json_exists`. This is the right kind of index to use for `json_exists`, because there are only two possible return values for a condition (true and false).

This is illustrated by [Example 26-1](#).

[Example 26-2](#) creates a bitmap index for a value returned by `json_value`. This is an appropriate index to use *if* there are only few possible values for field `CostCenter` in your data.

Example 26-1 Creating a Bitmap Index for `JSON_EXISTS`

```
CREATE BITMAP INDEX has_zipcode_idx
  ON j_purchaseorder (json_exists(po_document,
    '$.ShippingInstructions.Address.zipCode'));
```

Example 26-2 Creating a Bitmap Index for `JSON_VALUE`

```
CREATE BITMAP INDEX cost_ctr_idx
  ON j_purchaseorder (json_value(po_document, '$.CostCenter'));
```

26.4 Creating `JSON_VALUE` Function-Based Indexes

You can create a function-based index for SQL/JSON function `json_value`. You can use the standard syntax for this, explicitly specifying function `json_value`, or you can use the simple dot-notation syntax. Indexes created in either of these ways can be used with both dot-notation queries and `json_value` queries.

[Example 26-4](#) creates a function-based index for `json_value` on field `PONumber` of the object that is in column `po_document` of table `j_purchaseorder`. The object is passed as the path-expression context item.

The use of `ERROR ON ERROR` here means that if the data contains a record that has *no* `PONumber` field, has *more than one* `PONumber` field, or has a `PONumber` field with a *non-number* value then index creation fails. And if the index exists then trying to insert such a record fails.

An alternative is to create an index using the simplified syntax described in [Simple Dot-Notation Access to JSON Data](#). [Example 26-3](#) illustrates this; it indexes both scalar and non-scalar results, corresponding to what a dot-notation query can return.

The indexes created in both [Example 26-4](#) and [Example 26-3](#) can be picked up for either a query that uses dot-notation syntax or a query that uses `json_value`.

If the index of [Example 26-3](#) is picked up for a `json_value` query then filtering is applied after index pickup, to test for the correct field value. Non-scalar values can be stored in this index, since dot-notation queries can return such values, but a `json_value` query cannot, so such values are filtered out after index pickup.

If you want to allow indexing of data that might be missing the field targeted by a `json_value` expression, then use a `NULL ON EMPTY` clause, together with an `ERROR ON ERROR` clause. [Example 26-5](#) illustrates this.

Oracle *recommends* that you create a function-based index for `json_value` using one of these forms:

- **Dot-notation syntax**
The indexed values correspond to the flexible behavior of dot-notation queries, which return JSON values whenever possible. They can include non-scalar JSON values (JSON objects and arrays). They can match dot-notation queries in addition to `json_value` queries. The index is used to come up with an initial set of matches, which are then filtered according to the specifics of the query. For example, any indexed values that are not JSON scalars are filtered out.
- A `json_value` expression that specifies a **RETURNING** data type, uses **ERROR ON ERROR** (and optionally uses `NULL ON EMPTY`).
The indexed values are only (non-null) scalar values of the specified data type. The index can nevertheless be used in dot-notation queries that lead to such a scalar result.

Indexes created in either of these ways can thus be used with both dot-notation queries and `json_value` queries.

Example 26-3 Creating a Function-Based Index for a JSON Field: Dot Notation

```
CREATE UNIQUE INDEX po_num_idx2 ON j_purchaseorder po
(po.po_document.PONumber);
```

Example 26-4 Creating a Function-Based Index for a JSON Field: JSON_VALUE

```
CREATE UNIQUE INDEX po_num_idx1
ON j_purchaseorder (json_value(po_document, '$.PONumber'
RETURNING NUMBER ERROR ON ERROR));
```

Example 26-5 Specifying NULL ON EMPTY for a JSON_VALUE Function-Based Index

Because of clause `NULL ON EMPTY`, index `po_ref_idx1` can index JSON documents that have no Reference field.

```
CREATE UNIQUE INDEX po_ref_idx1
ON j_purchaseorder (json_value(po_document, '$.Reference'
```

```
RETURNING VARCHAR2(200) ERROR ON ERROR
NULL ON EMPTY));
```

Related Topics

- [Empty-Field Clause for SQL/JSON Query Functions](#)
SQL/JSON query functions `json_value`, `json_query`, and `json_table` accept an optional `ON EMPTY` clause, which specifies the handling to use when a targeted JSON field is absent from the data queried. This clause and the default behavior (no `ON EMPTY` clause) are described here.
- [Using GeoJSON Geographic Data](#)
GeoJSON objects are JSON objects that represent geographic data. Examples are provided of creating GeoJSON data, indexing it, and querying it.

26.5 Using a JSON_VALUE Function-Based Index with JSON_TABLE Queries

An index created using `json_value` with `ERROR ON ERROR` can be used for a query involving `json_table`, if the `WHERE` clause refers to a column projected by `json_table`, and the effective SQL/JSON path that targets that column matches the indexed path expression.

The index acts as a constraint on the indexed path, to ensure that only one (non-null) scalar JSON value is projected for each item in the JSON collection.

The query in [Example 26-6](#) thus makes use of the index created in [Example 26-4](#).

Note:

A function-based index created using a `json_value` expression or dot notation can be picked up for a corresponding occurrence in a query `WHERE` clause only if the occurrence is used in a SQL *comparison* condition, such as `>=`. In particular, it is not picked up for an occurrence used in condition `IS NULL` or `IS NOT NULL`.

See *Oracle Database SQL Language Reference* for information about SQL comparison conditions.

Example 26-6 Use of a JSON_VALUE Function-Based Index with a JSON_TABLE Query

```
SELECT jt.*
FROM j_purchaseorder po,
     json_table(po.po_document, '$'
               COLUMNS po_number NUMBER(5) PATH '$.PONumber',
                        reference  VARCHAR2(30 CHAR) PATH '$.Reference',
                        requestor  VARCHAR2(32 CHAR) PATH '$.Requestor',
                        userid     VARCHAR2(10 CHAR) PATH '$.User',
                        costcenter VARCHAR2(16 CHAR) PATH '$.CostCenter') jt
WHERE po_number = 1600;
```

26.6 Using a JSON_VALUE Function-Based Index with JSON_EXISTS Queries

An index created using SQL/JSON function `json_value` with `ERROR ON ERROR` can be used for a query involving SQL/JSON condition `json_exists`, provided the query path expression has a filter expression that contains only a *path-expression comparison* or multiple such comparisons separated by `&&`.

In order for a `json_value` function-based index to be picked up for one of the comparisons of the query, the type of that comparison must be the same as the returning SQL data type for the index. The SQL data types used are those mentioned for item methods `double()`, `number()`, `timestamp()`, `date()`, and `string()` — see [SQL/JSON Path Expression Item Methods](#).

For example, if the index returns a number then the comparison type must also be number. If the query filter expression contains more than one comparison that matches a `json_value` index, the optimizer chooses one of the indexes.

The *type of a comparison* is determined as follows:

1. If the SQL data types of the two comparison terms (sides of the comparison) are different then the type of the comparison is *unknown*, and the index is not picked up. Otherwise, the types are the same, and this type is the type of the comparison.
2. If a comparison term is of SQL data type *string* (a text literal) then the type of the comparison is the *type of the other comparison term*.
3. If a comparison term is a *path expression* with a function step whose *item method imposes a SQL match type* then that is also the type of that comparison term. The item methods that impose a SQL match type are `double()`, `number()`, `timestamp()`, `date()`, and `string()`.
4. If a comparison term is a *path expression* with *no* such function step then its type is SQL *string* (text literal).

Example 26-4 creates a function-based index for `json_value` on field `PONumber`. The index return type is `NUMBER`.

Each of the queries [Example 26-7](#), [Example 26-8](#), and [Example 26-9](#) can make use of this index when evaluating its `json_exists` condition. Each of these queries uses a comparison that involves a simple path expression that is relative to the absolute path expression `$.PONumber`. The relative simple path expression in each case targets the current filter item, `@`, but in the case of [Example 26-9](#) it transforms (casts) the matching data to SQL data type `NUMBER`.

Example 26-7 JSON_EXISTS Query Targeting Field Compared to Literal Number

This query makes use of the index because:

1. One comparison term is a path expression with no function step, so its type is SQL *string* (text literal).
2. Because one comparison term is of type *string*, the comparison has the type of the other term, which is *number* (the other term is a numeral).

3. The type of the (lone) comparison is the same as the type returned by the index: *number*.

```
SELECT count(*) FROM j_purchaseorder
WHERE json_exists(po_document, '$.PONumber?(@ > 1500)');
```

Example 26-8 JSON_EXISTS Query Targeting Field Compared to Variable Value

This query can make use of the index because:

1. One comparison term is a path expression with no function step, so its type is SQL *string* (text literal).
2. Because one comparison term is of type string, the comparison has the type of the other term, which is *number* (the other term is a variable that is bound to a number).
3. The type of the (lone) comparison is the same as the type returned by the index: *number*.

```
SELECT count(*) FROM j_purchaseorder
WHERE json_exists(po_document, '$.PONumber?(@ > $d)'
                 PASSING 1500 AS "d");
```

Example 26-9 JSON_EXISTS Query Targeting Field Cast to Number Compared to Variable Value

This query can make use of the index because:

1. One comparison term is a path expression with a function step whose item method (`number()`) transforms the matching data to a *number*, so the type of that comparison term is SQL *number*.
2. The other comparison term is a numeral, which has SQL type *number*. The types of the comparison terms match, so the comparison has this same type, *number*.
3. The type of the (lone) comparison is the same as the type returned by the index: *number*.

```
SELECT count(*) FROM j_purchaseorder
WHERE json_exists(po_document, '$.PONumber?(@.number() > $d)'
                 PASSING 1500 AS "d");
```

Example 26-10 JSON_EXISTS Query Targeting a Conjunction of Field Comparisons

Just as for [Example 26-7](#), this query can make use of the index on field `PONumber`. If a `json_value` index is also defined for field `Reference` then the optimizer chooses which index to use for this query.

```
SELECT count(*) FROM j_purchaseorder
WHERE json_exists(po_document, '$?(@.PONumber > 1500
                                && @.Reference ==
                                "ABULL-20140421")');
```

Related Topics

- [Creating JSON_VALUE Function-Based Indexes](#)
You can create a function-based index for SQL/JSON function `json_value`. You can use the standard syntax for this, explicitly specifying function `json_value`, or you can use the simple dot-notation syntax. Indexes created in either of these ways can be used with both dot-notation queries and `json_value` queries.
- [SQL/JSON Path Expressions](#)
Oracle Database provides SQL access to JSON data using SQL/JSON path expressions.

26.7 Data Type Considerations for JSON_VALUE Indexing and Querying

By default, SQL/JSON function `json_value` returns a `VARCHAR2` value. When you create a function-based index using `json_value`, unless you use a `RETURNING` clause to specify a different return data type, the index is not picked up for a query that expects a non-`VARCHAR2` value.

For example, in the query of [Example 26-11](#), `json_value` uses `RETURNING NUMBER`. The index created in [Example 26-4](#) can be picked up for this query, because the indexed `json_value` expression specifies a return type of `NUMBER`.

But the index created in [Example 26-3](#) does not use `RETURNING NUMBER` (the return type is `VARCHAR2(4000)`, by default), so it cannot be picked up for a such a query.

Now consider the queries in [Example 26-12](#) and [Example 26-13](#), which use `json_value` without a `RETURNING` clause, so that the value returned is of type `VARCHAR2`.

In [Example 26-12](#), SQL function `to_number` explicitly converts the `VARCHAR2` value returned by `json_value` to a number. Similarly, in [Example 26-13](#), comparison condition `>` (greater-than) implicitly converts the value to a number.

Neither of the indexes of [Example 26-4](#) and [Example 26-3](#) is picked up for either of these queries. The queries might return the right results in each case, because of type-casting, but the indexes cannot be used to evaluate the queries.

Consider also what happens if some of the data cannot be converted to a particular data type. For example, given the queries in [Example 26-11](#), [Example 26-12](#), and [Example 26-13](#), what happens to a `PONumber` value such as "alpha"?

For [Example 26-12](#) and [Example 26-13](#), the query stops in error because of the attempt to cast the value to a number. For [Example 26-11](#), however, because the default error handling behavior is `NULL ON ERROR`, the non-number value "alpha" is simply filtered out. The value is indexed, but it is ignored for the query.

Similarly, if the query used, say, `DEFAULT '1000' ON ERROR`, that is, if it specified a numeric default value, then no error would be raised for the value "alpha": the default value of 1000 would be used.

Example 26-11 JSON_VALUE Query with Explicit RETURNING NUMBER

```
SELECT count(*) FROM j_purchaseorder po
       WHERE json_value(po_document, '$.PONumber' RETURNING NUMBER) > 1500;
```

Example 26-12 JSON_VALUE Query with Explicit Numerical Conversion

```
SELECT count(*) FROM j_purchaseorder po
WHERE to_number(json_value(po_document, '$.PONumber')) > 1500;
```

Example 26-13 JSON_VALUE Query with Implicit Numerical Conversion

```
SELECT count(*) FROM j_purchaseorder po
WHERE json_value(po_document, '$.PONumber') > 1500;
```

26.8 Indexing Multiple JSON Fields Using a Composite B-Tree Index

To index multiple fields of a JSON object, you first create virtual columns for them. Then you create a composite B-tree index on the virtual columns.

[Example 26-14](#) and [Example 26-15](#) illustrate this. [Example 26-14](#) creates virtual columns `userid` and `costcenter` for JSON object fields `User` and `CostCenter`, respectively.

[Example 26-15](#) creates a composite B-tree index on the virtual columns of [Example 26-14](#).

A SQL query that references either the virtual columns or the corresponding JSON data (object fields) picks up the composite index. This is the case for both of the queries in [Example 26-16](#).

These two queries have the same effect, including the same performance. However, the first query form does not target the JSON data itself; it targets the virtual columns that are used to index that data.

The data does not depend logically on any indexes implemented to improve query performance. If you want this independence from implementation to be reflected in your code, then use the second query form. Doing that ensures that the query behaves the same functionally with or without the index — the index serves only to improve performance.

Example 26-14 Creating Virtual Columns For JSON Object Fields

```
ALTER TABLE j_purchaseorder ADD (userid VARCHAR2(20)
GENERATED ALWAYS AS (json_value(po_document, '$.User' RETURNING
VARCHAR2(20))));
```

```
ALTER TABLE j_purchaseorder ADD (costcenter VARCHAR2(6)
GENERATED ALWAYS AS (json_value(po_document, '$.CostCenter'
RETURNING VARCHAR2(6))));
```

Example 26-15 Creating a Composite B-tree Index For JSON Object Fields

```
CREATE INDEX user_cost_ctr_idx on j_purchaseorder(userid, costcenter);
```

Example 26-16 Two Ways to Query JSON Data Indexed With a Composite Index

```

SELECT po_document FROM j_purchaseorder WHERE userid      = 'ABULL'
                                             AND costcenter = 'A50';

SELECT po_document
FROM j_purchaseorder WHERE json_value(po_document, '$.User') =
'ABULL'
                           AND json_value(po_document, '$.CostCenter') =
'A50';

```

26.9 JSON Search Index: Ad Hoc Queries and Full-Text Search

A JSON search index is a *general* index. It can improve the performance of both (1) ad hoc structural queries, that is, queries that you might not anticipate or use regularly, and (2) full-text search. It is an Oracle Text index that is designed specifically for use with JSON data.

**Note:**

If you created a JSON search index using Oracle Database 12c Release 1 (12.1.0.2) then Oracle recommends that you *drop* that index and *create a new search index* for use with later releases, using `CREATE SEARCH INDEX` as shown here.

Introduction to JSON Search Indexes

You create a JSON search index using `CREATE SEARCH INDEX` with the keywords `FOR JSON`. [Example 26-17](#) illustrates this.

If the name of your JSON search index is present in the execution plan for your query, then you know that the index was in fact picked up for that query. You will see a line similar to that shown in [Example 26-18](#).

You can specify a `PARAMETERS` clause to override the default settings of certain configurable options. By default (no `PARAMETERS` clause), the index is synchronized on commit and both text and numeric ranges are indexed.

A JSON search index is maintained asynchronously, on demand. You can thus defer the cost of index maintenance, performing it at commit time only or at some time when database load is reduced. This can improve DML performance. It can also improve index maintenance performance by enabling bulk loading of unsynchronized index rows when an index is synchronized. On the other hand, asynchronous maintenance of an index means that until it is synchronized the index is not used for data that has been modified or newly inserted.

If your queries that make use of a JSON search index never involve numeric ranges then you can save some index maintenance time and some disk space by specifying `TEXT` for parameter `SEARCH_ON`. The default value of `SEARCH_ON` is `TEXT_VALUE`, which means index numeric ranges as well as text.

 **Note:**

To alter a JSON search index `j_s_idx`, you use `ALTER INDEX j_s_idx REBUILD ...` (*not* `ALTER SEARCH INDEX j_s_idx ...`).

Example 26-17 Creating a JSON Search Index

```
CREATE SEARCH INDEX po_search_idx ON j_purchaseorder (po_document) FOR
JSON;
```

Example 26-18 Execution Plan Indication that a JSON Search Index Is Used

```
|* 2| DOMAIN INDEX | PO_SEARCH_IDX | | | 4 (0)
```

Full-Text Search of JSON Data

You can use SQL/JSON condition `json_textcontains` in a `CASE` expression or the `WHERE` clause of a `SELECT` statement to perform a full-text search of JSON data that is stored in a `VARCHAR2`, `BLOB`, or `CLOB` column.

Oracle Text technology underlies SQL condition `json_textcontains`. This means, for instance, that you can query for text that is near some other text, or query use fuzzy pattern-matching.

To be able to use condition `json_textcontains`, you first must create a JSON search index. If you do not, then an error is raised when `json_textcontains` is used.

[Example 26-19](#) shows a full-text query that finds purchase-order documents that contain the keyword `Magic` in any of the line-item part descriptions.

 **Note:**

For powerful full-text search, use Oracle SQL function `json_textcontains`, which requires that you create a JSON search index. As a less-powerful alternative, if you do not create a JSON search index, and you just want simple string pattern-matching in a filter condition, you can use any of the pattern-matching comparisons: `has substring`, `starts with`, `like`, `like_regex`, or `eq_regex`.

Example 26-19 Full-Text Query of JSON Data

```
SELECT po_document FROM j_purchaseorder
WHERE json_textcontains(po_document, '$.LineItems.Part.Description',
'Magic');
```


Example 26-20 Full-Text Query of JSON Data, with Escaped Search Pattern

If the third argument to condition `json_textcontains` contains a character or a word that is *reserved* with respect to Oracle Text search then you must *escape* that character or word.

The queries here search for documents where the part description contains both `Magic` and `Christmas`. They use the Oracle Text query operators `&` and `and`, which are reserved with respect to Oracle Text search.

The first query escapes the single character `&`, using backslash (`\`). The second escapes the entire search pattern, using braces (`{, }`), because it contains the reserved word `and`. (It could alternatively use `{and}` to escape only the characters in operator `and`).

```
SELECT po_document FROM j_purchaseorder
       WHERE json_textcontains(po_document, '$.LineItems.Part.Description',
                              'Magic \& Christmas');
```

```
SELECT po_document FROM j_purchaseorder
       WHERE json_textcontains(po_document, '$.LineItems.Part.Description',
                              '{Magic and Christmas}');
```

Ad Hoc Queries of JSON Data

[Example 26-21](#) shows some *non* full-text queries of JSON data that also make use of the JSON search index created in [Example 26-17](#).

Example 26-21 Some Ad Hoc JSON Queries

This query selects documents that contain a shipping instructions address that includes a country.

```
SELECT po_document FROM j_purchaseorder
       WHERE json_exists(po_document,
                        '$.ShippingInstructions.Address.country');
```

This query selects documents that contain user `AKH00` where there are more than 8 items ordered. It takes advantage of numeric-range indexing.

```
SELECT po_document FROM j_purchaseorder
       WHERE json_exists(po_document, '$?(@.User == "AKH00"
                                     && @.LineItems.Quantity > 8)');
```

This query selects documents where the user is `AKH00`. It uses `json_value` instead of `json_exists` in the `WHERE` clause.

```
SELECT po_document FROM j_purchaseorder
       WHERE json_value(po_document, '$.User') = 'ABULL';
```

Related Topics

- [Indexes for JSON Data](#)
You can index JSON data as you would any data of the type you use to store it. In addition, you can define a JSON search index, which is useful for both ad hoc structural queries and full-text queries.
- [JSON Data Guide](#)
A JSON data guide lets you discover information about the structure and content of JSON documents stored in Oracle Database.

See Also:

- *Oracle Database SQL Language Reference* for information about condition `json_textcontains`
- *Oracle Text Reference* for information about the `PARAMETERS` clause for `CREATE SEARCH INDEX`
- *Oracle Text Reference* for information about the `PARAMETERS` clause for `ALTER INDEX ... REBUILD`
- *Oracle Text Reference* for information about synchronizing a JSON search index
- *Oracle Text Application Developer's Guide* for guidance about optimizing and tuning the performance of a JSON search index
- *Oracle Text Reference* for information about the words and characters that are reserved with respect to Oracle Text search, and *Oracle Text Reference* for information about how to escape them.

In-Memory JSON Data

A column of JSON data can be stored in the In-Memory Column Store (IM column store) to improve query performance.

- [Overview of In-Memory JSON Data](#)
You can move a table with a column of JSON data to the In-Memory Column Store (IM column store), to improve the performance of queries that share costly expressions by caching the expression results. This is especially useful for analytical queries that scan a large number of small JSON documents.
- [Populating JSON Data Into the In-Memory Column Store](#)
You use `ALTER TABLE INMEMORY` to populate a table with a column of JSON data into the In-Memory Column Store (IM column store), to improve the performance of queries that share costly expressions by caching the results of evaluating those expressions.
- [Upgrading Tables With JSON Data For Use With the In-Memory Column Store](#)
A table with JSON columns created using a database that did not have a compatibility setting of at least 12.2 or did not have `max_string_size = extended` must first be upgraded, before it can be populated into the In-Memory Column Store (IM column store). To do this, run script `rdbms/admin/utlimcjson.sql`.

 **See Also:**

Oracle Database In-Memory Guide

27.1 Overview of In-Memory JSON Data

You can move a table with a column of JSON data to the In-Memory Column Store (IM column store), to improve the performance of queries that share costly expressions by caching the expression results. This is especially useful for analytical queries that scan a large number of small JSON documents.

The IM column store is supported only for JSON documents smaller than 32,767 bytes. If you have a mixture of document sizes, those documents that are larger than 32,767 bytes are processed without the In-Memory optimization. For better performance, consider breaking up documents larger than 32,767 bytes into smaller documents.

The IM column store is an optional SGA pool that stores copies of tables and partitions in a special columnar format optimized for rapid scans. The IM column store supplements the row-based storage in the database buffer cache. You do not need to load the same object into both the IM column store and the buffer cache. The two caches are kept transactionally consistent. The database transparently sends online transaction processing (OLTP) queries (such as primary-key lookups) to the buffer cache and analytic and reporting queries to the IM column store.

You can think of the use of JSON data in memory as improving the performance of SQL/JSON path access. SQL/JSON functions `json_table`, `json_query`, and `json_value`, and SQL condition `json_exists` all accept a SQL/JSON path argument, and they can all benefit from loading JSON data into the IM column store. (Full-text search using SQL/JSON function `json_textcontains` does *not* benefit from the IM column store.) Once JSON documents have been loaded into memory, any subsequent path-based operations on them use the In-Memory representation, which avoids the overhead associated with reading and parsing the on-disk format.

If queried JSON data is populated into the IM column store, and if there are function-based indexes that can apply to that data, the optimizer chooses whether to use an index or to scan the data in memory. In general, if index probing results in few documents then a functional index can be preferred by the optimizer. In practice this means that the optimizer can prefer a functional index for very selective queries or DML statements.

On the other hand, if index probing results in many documents then the optimizer might choose to scan the data in memory, by scanning the function-based index expression as a virtual-column expression.

Ad hoc queries, that is, queries that are not used frequently to target a given SQL/JSON path expression, benefit in a general way from populating JSON data into the IM column store, by quickly scanning the data. But if you have some frequently used queries then you can often further improve their performance in these ways:

- Creating *virtual columns* that project scalar values (not under an array) from a column of JSON data and loading those virtual columns into the IM column store.
- Creating a *materialized view* on a frequently queried `json_table` expression and loading the view into the IM column store.

However, if you have a function-based index that projects a scalar value using function `json_value` then you need not explicitly create a virtual column to project it. As mentioned above, in this case the function-based index expression is automatically loaded into the IM column store as a virtual column. The optimizer can choose, based on estimated cost, whether to scan the function-based index in the usual manner or to scan the index expression as a virtual-column expression.

 **Note:**

- The advantages of a virtual column over a materialized view are that you can build an index on it and you can obtain statistics on it for the optimizer.
- Virtual columns, like columns in general, are subject to the 1000-column limit for a given table.

Prerequisites For Using JSON Data In Memory

To be able to take advantage of the IM column store for JSON data, the following must *all* be true:

- Database compatibility is 12.2.0.0 or higher.
- The values set for `max_string_size` in the Oracle instance startup configuration file must be 'extended'.

- Sufficient SGA memory must be configured for the IM column store.
- A DBA has specified that the tablespace, table, or materialized view that contains the JSON columns is eligible for population into the IM column store, using keyword `INMEMORY` in a `CREATE` or `ALTER` statement.
- Initialization parameters are set as follows:
 - `INMEMORY_EXPRESSIONS_USAGE` is `STATIC_ONLY` or `ENABLE`.
`ENABLE` allows In-Memory materialization of dynamic expressions, if used in conjunction with PL/SQL procedure `DBMS_INMEMORY.ime_capture_expressions`.
 - `INMEMORY_VIRTUAL_COLUMNS` is `ENABLE`, meaning that the IM column store populates all virtual columns. (The default value is `MANUAL`.)
- The columns storing the JSON data must each have `is json` check constraints. (That is, the data must be known to be JSON data.)

You can check the value of each initialization parameter using command `SHOW PARAMETER`. (You must be logged in as database user `SYS` or equivalent for this.) For example:

```
SHOW PARAMETER INMEMORY_VIRTUAL_COLUMNS
```

Related Topics

- [Populating JSON Data Into the In-Memory Column Store](#)
 You use `ALTER TABLE INMEMORY` to populate a table with a column of JSON data into the In-Memory Column Store (IM column store), to improve the performance of queries that share costly expressions by caching the results of evaluating those expressions.

27.2 Populating JSON Data Into the In-Memory Column Store

You use `ALTER TABLE INMEMORY` to populate a table with a column of JSON data into the In-Memory Column Store (IM column store), to improve the performance of queries that share costly expressions by caching the results of evaluating those expressions.

The IM column store is an optional SGA pool that stores copies of tables and partitions in a special columnar format optimized for rapid scans. The IM column store supplements the row-based storage in the database buffer cache. (It does not replace the buffer cache, but you do not need to load the same object into both the IM column store and the buffer cache. The two caches are kept transactionally consistent.)

You specify that a table with a given JSON column (that is, a column that has an `is json` check constraint) is to be populated into the IM column store by marking the table as `INMEMORY`. [Example 27-1](#) illustrates this.

The IM column store is used for queries of documents that are smaller than 32,767 bytes. Queries of documents that are larger than that do not benefit from the IM column store.

 **Note:**

If a JSON column in a table that is to be populated into the IM column store was created using a database that did not have a compatibility setting of at least 12.2 or did not have `max_string_size` set to `extended` (this is the case prior to Oracle Database 12c Release 2 (12.2.0.1), for instance) then you must first run script `rdbms/admin/utlimcjson.sql`. It prepares *all* existing tables that have JSON columns to take advantage of the In-Memory JSON processing that was added in Release 12.2.0.1. See [Upgrading Tables With JSON Data For Use With the In-Memory Column Store](#).

After you have marked a table that has JSON columns as `INMEMORY`, an *In-Memory virtual column* is added to it for each JSON column. The corresponding virtual column is used for queries of a given JSON column. The virtual column contains the same JSON data as the corresponding JSON column, but in an Oracle binary format, **OSON**.

Example 27-1 Populating JSON Data Into the IM Column Store

```
SELECT COUNT(1) FROM j_purchaseorder
  WHERE json_exists(po_document,
                   '$.ShippingInstructions?(@.Address.zipCode ==
99236)');

-- The execution plan shows: TABLE ACCESS FULL

-- Specify table as INMEMORY, with default PRIORITY setting of NONE,
-- so it is populated only when a full scan is triggered.

ALTER TABLE j_purchaseorder INMEMORY;

-- Query the table again, to populate it into the IM column store.
SELECT COUNT(1) FROM j_purchaseorder
  WHERE json_exists(po_document,
                   '$.ShippingInstructions?(@.Address.zipCode ==
99236)');

-- The execution plan for the query now shows: TABLE ACCESS INMEMORY
FULL
```

 **See Also:**

Oracle Database In-Memory Guide

27.3 Upgrading Tables With JSON Data For Use With the In-Memory Column Store

A table with JSON columns created using a database that did not have a compatibility setting of at least 12.2 or did not have `max_string_size = extended` must first be upgraded, before it can be populated into the In-Memory Column Store (IM column store). To do this, run script `rdbms/admin/utlimcjson.sql`.

Script `rdbms/admin/utlimcjson.sql` upgrades *all* existing tables that have JSON columns so they can be populated into the IM column store. To use it, *all* of the following must be true:

- Database parameter `compatible` must be set to 12.2.0.0 or higher.
- Database parameter `max_string_size` must be set to `extended`.
- The JSON columns being upgraded must have an `is json` check constraint defined on them.

Related Topics

- [Overview of In-Memory JSON Data](#)
You can move a table with a column of JSON data to the In-Memory Column Store (IM column store), to improve the performance of queries that share costly expressions by caching the expression results. This is especially useful for analytical queries that scan a large number of small JSON documents.

A

Oracle Database JSON Restrictions

The restrictions associated with Oracle support of JSON data in Oracle Database are listed here.

Unless otherwise specified, an error is raised if a specified limitation is not respected.

- General
 - *Number of nesting levels for a JSON object or array:* 1000, maximum.
 - *JSON field name length:* 32767 bytes, maximum.
- SQL/JSON functions
 - *Return-value length:* 32767 bytes, maximum.
 - *Path length:* 4K bytes, maximum.
 - *Number of path steps:* 65535, maximum.
- Simplified JSON syntax
 - *Path length:* 4K bytes, maximum.
 - *Path component length:* 128 bytes, maximum.
- JSON search index
 - *Field name length:* 255 bytes, maximum. If a document has a field name longer than 255 bytes then it might not be completely indexed. In that case, an error is recorded in database view `CTX_USER_INDEX_ERRORS`.

Note:

You **must rebuild** any JSON search indexes and Oracle Text indexes created prior to Oracle Database 18c if they index JSON data that contains object fields with names longer than 64 bytes. See *Oracle Database Upgrade Guide* for more information.

See Also:

Oracle Text Reference

- JSON data guide
 - *Path length:* 4000 bytes, maximum. A path longer than 4000 bytes is *ignored* by a data guide.
 - *Number of children under a parent node:* 1024, maximum. A node that has more than 1024 children is *ignored* by a data guide.

-
- *Field value length*: 32767 bytes. If a JSON field has a value longer than 32767 bytes then the data guide reports the length as 32767.
 - *Zero-length field name*: A zero-length (empty) object field name ("") is not supported for use with JSON data guide. Data-guide behavior is undefined for JSON data that contains such a name.
 - OSON
 - *All field names, total size*: 65535 bytes, maximum.
 - *Field name length (each)*: 255 bytes, maximum.
 - *No duplicate fields*: If a JSON object with duplicate field names is represented using OSON then only one of these fields is present (kept).
 - PL/SQL getter method `JSON_OBJECT_T.get_keys()` returns at most 32767 field names for a given JSON object. An error is raised if it is applied to an object with more than 32767 fields.

B

Diagrams for Basic SQL/JSON Path Expression Syntax

Syntax diagrams and corresponding Backus-Naur Form (BNF) syntax descriptions are presented for the basic SQL/JSON path expression syntax.

The basic syntax of SQL/JSON path expression is explained in [Basic SQL/JSON Path Expression Syntax](#). This topic recapitulates that information in the form of syntax diagrams and BNF descriptions.

Figure B-1 json_basic_path_expression




Figure B-2 json_absolute_path_expression




Figure B-3 json_nonfunction_steps




Figure B-4 json_object_step




Figure B-5 json_field_name





Figure B-6 json_array_step



 **Note:**

- Array indexing is zero-based, so *integer* is a non-negative integer (0, 1, 2, 3,...).
- The first *integer* of a range (*integer to integer*) must be less than the second.
- The array elements must be specified by indexes in ascending order, without repetitions.

A compile-time error is raised if any of these syntax rules is violated.

Figure B-7 json_function_step




Figure B-8 json_item_method




Figure B-9 json_filter_expr




Figure B-10 json_cond

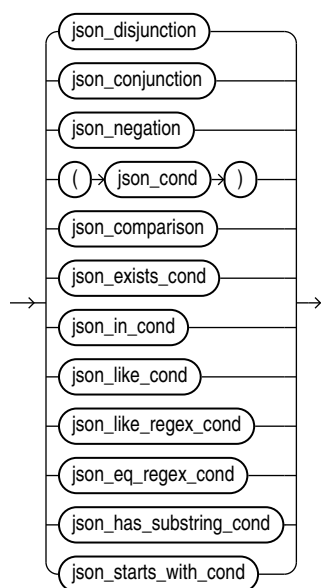


Figure B-11 json_conjunction




Figure B-12 json_comparison




Figure B-13 json_relative_path-expr







Figure B-14 json_compare_pred**Figure B-15** json_var**Figure B-16** json_scalar**Note:**

json_number is a JSON number: a decimal numeral, possibly signed and possibly including a decimal exponent.

See Also:

Oracle Database SQL Language Reference for information about Oracle syntax diagrams

Index

Symbols

! filter predicate, SQL/JSON path expressions, [13-2](#)

!= comparison filter predicate, SQL/JSON path expressions, [13-2](#)

&& filter predicate, SQL/JSON path expressions, [13-2](#)

< comparison filter predicate, SQL/JSON path expressions, [13-2](#)

<= comparison filter predicate, SQL/JSON path expressions, [13-2](#)

<> comparison filter predicate, SQL/JSON path expressions, [13-2](#)

== comparison filter predicate, SQL/JSON path expressions, [13-2](#)

> comparison filter predicate, SQL/JSON path expressions, [13-2](#)

>= comparison filter predicate, SQL/JSON path expressions, [13-2](#)

|| filter predicate, SQL/JSON path expressions, [13-2](#)

\$, SQL/JSON path expressions
for a SQL/JSON variable, [13-2](#)
for the context item, [13-2](#)

A

abs() item method, SQL/JSON path expressions, [13-11](#)

ABSENT ON NULL, SQL/JSON generation functions, [21-1](#)

absolute path expression, [13-2](#)
syntax, [B-1](#)

ad hoc search of JSON data, [26-10](#)

add_vc trigger procedure, [20-30](#)

add_virtual_columns, DBMS_JSON PL/SQL procedure, [20-23](#), [20-24](#), [20-27](#)

adding virtual columns for JSON fields, [20-23](#)
based on a data guide-enabled search index, [20-27](#)
based on a hierarchical data guide, [20-24](#)

ALL_JSON_COLUMNS view, [4-3](#)

ALL_JSON_DATAGUIDE_FIELDS view, [20-12](#)

ALL_JSON_DATAGUIDES view, [20-12](#)

array element, JSON, [2-2](#)

array index, [13-2](#)

array range specification, [13-2](#)

array step, SQL/JSON path expressions, [13-2](#)
syntax, [B-1](#)

array, JSON, [2-2](#)

ASCII keyword, SQL functions, [14-1](#)

B

basic SQL/JSON path expression, [13-2](#)
BNF description, [B-1](#)
diagrams, [B-1](#)

BLOB storage, using for JSON data, [3-1](#)

BNF syntax descriptions, basic SQL/JSON path expression, [B-1](#)

boolean() item method, SQL/JSON path expressions, [13-11](#)

booleanOnly() item method, SQL/JSON path expressions, [13-11](#)

C

canonical form of a JSON number, [14-1](#)

case-sensitivity
in data-guide field
o:preferred_column_name, [20-9](#)
in query dot notation, [12-1](#)
in SQL/JSON path expression, [13-2](#)
JSON and SQL, [xvi](#)
strict and lax JSON syntax, [5-2](#)

ceiling() item method, SQL/JSON path expressions, [13-11](#)

change trigger, data guide, [20-30](#)
user-defined, [20-32](#)

character sets, [6-1](#)

check constraint used to ensure well-formed JSON data, [4-1](#)

child COLUMNS clause, json_table SQL/JSON function, [18-5](#)

client, using to retrieve JSON Data, [3-1](#)

CLOB storage, using for JSON data, [3-1](#)

COLUMNS clause, json_table SQL/JSON function, [18-5](#)

columns of JSON data, [1-2](#)

compare predicate, SQL/JSON path expressions
syntax, [B-1](#)

comparison filter predicates, SQL/JSON path
expressions, [13-2](#)

comparison in SQL/JSON path expression,
types, [13-16](#)

comparison, SQL/JSON path expressions
syntax, [B-1](#)

compatibility of data types, item methods, [13-11](#)

condition, SQL/JSON path expressions
syntax, [B-1](#)

conditions, Oracle SQL
json_equal, [1](#)

conditions, SQL/JSON
is json, [5-1](#)
and JSON null, [2-2](#)
is not json, [5-1](#)
and JSON null, [2-2](#)
json_exists, [15-1](#)
indexing, [26-3](#)
json_textcontains, [26-10](#)

conjunction, SQL/JSON path expressions
syntax, [B-1](#)

context item, SQL/JSON path expressions, [13-2](#)

create_view_on_path, DBMS_JSON PL/SQL
procedure, [20-15](#), [20-19](#)

create_view, DBMS_JSON PL/SQL procedure,
[20-15](#), [20-16](#)

D

data guide
change trigger, [20-30](#)
user-defined, [20-32](#)
fields, [20-9](#)
flat, [20-41](#)
hierarchical, [20-47](#)
multiple for the same JSON column, [20-35](#)
overview, [20-2](#)

data types for JSON columns, [3-1](#)

date and time formats, ISO 8601, [13-14](#)

date() item method, SQL/JSON path
expressions, [13-11](#)

DBA_JSON_COLUMNS view, [4-3](#)

DBA_JSON_DATAGUIDE_FIELDS view, [20-12](#)

DBA_JSON_DATAGUIDES view, [20-12](#)

DBMS_JSON.add_virtual_columns PL/SQL
procedure, [20-23](#), [20-24](#), [20-27](#)

DBMS_JSON.create_view PL/SQL procedure,
[20-15](#), [20-16](#)

DBMS_JSON.create_view_on_path PL/SQL
procedure, [20-15](#), [20-19](#)

DBMS_JSON.drop_virtual_columns PL/SQL
procedure, [20-23](#), [20-30](#)

DBMS_JSON.FORMAT_FLAT, [20-7](#), [20-9](#), [20-41](#)

DBMS_JSON.FORMAT_HIERARCHICAL, [20-7](#),
[20-16](#), [20-24](#), [20-47](#)

DBMS_JSON.get_index_dataguide PL/SQL
function, [20-7](#), [20-9](#), [20-16](#), [20-24](#)

DBMS_JSON.PRETTY, [20-16](#), [20-41](#), [20-47](#)

DBMS_JSON.rename_column PL/SQL
procedure, [20-9](#)

descendant step, SQL/JSON path expressions,
[13-2](#)

diagrams, basic SQL/JSON path expression
syntax, [B-1](#)

disjunction, SQL/JSON path expressions
syntax, [B-1](#)

Document Object Model (DOM), [22-1](#)

DOM-like manipulation of JSON data, [22-1](#)

dot-notation access to JSON data, [12-1](#)
use with json_table SQL/JSON function, [18-1](#)

double() item method, SQL/JSON path
expressions, [13-11](#)

drop_virtual_columns, DBMS_JSON PL/SQL
procedure, [20-23](#), [20-30](#)

dropping virtual columns for JSON fields, [20-23](#),
[20-30](#)

duplicate field names in JSON objects, [5-2](#)

E

element of a JSON array, [2-2](#)

eq_regex filter predicate, SQL/JSON path
expressions, [13-2](#)

error clause, SQL query functions and conditions,
[14-5](#)

ERROR ON MISMATCH clause, json_value,
[14-9](#)

exists filter predicate, SQL/JSON path
expressions, [13-2](#)

EXISTS keyword, json_table SQL/JSON
function, [18-1](#)

EXTRA DATA clause, ON MISMATCH clause,
json_value, [14-9](#)

F

field name, SQL/JSON path expressions
syntax, [B-1](#)

field, JSON object, [2-2](#)

filter condition, SQL/JSON path expressions,
[13-2](#)

filter expression, SQL/JSON path expressions,
[13-2](#)

filter, SQL/JSON path expressions
syntax, [B-1](#)

floor() item method, SQL/JSON path
expressions, [13-11](#)

FOR ORDINALITY keywords, json_table SQL/JSON function, [18-1](#)

FORMAT JSON keywords
 json_table SQL/JSON function, [18-5](#)
 SQL/JSON generation functions, [21-1](#), [21-3](#)

FORMAT_FLAT, package DBMS_JSON, [20-7](#), [20-9](#), [20-41](#)

FORMAT_HIERARCHICAL, package
 DBMS_JSON, [20-7](#), [20-16](#), [20-24](#), [20-47](#)

full-text search of JSON data, [26-10](#)

function step, SQL/JSON path expressions, [13-2](#)
 syntax, [B-1](#)

functions, Oracle SQL
 json_serialize, [19-1](#)

functions, SQL
 json_dataguide, [20-7](#), [20-9](#)
 as an aggregate function, [20-35](#)
 hierarchical format, [20-47](#)
 pretty-print format, [20-47](#)
 json_mergepatch, [9-1](#)

functions, SQL/JSON
 json_array, [21-10](#)
 json_arrayagg, [21-13](#)
 json_object, [21-6](#)
 json_objectagg, [21-12](#)
 json_query, [17-1](#)
 json_table, [18-1](#)
 json_value, [16-1](#)
 function-based indexing, [26-3](#)
 indexing for geographic data, [24-1](#)
 null JSON value, [16-4](#)
 returning an object-type instance, [16-4](#)

G

generation of JSON data using SQL, [21-1](#)
 input SQL values, [21-3](#)

geographic JSON data, [24-1](#)

GeoJSON, [24-1](#)

geometric features in JSON, [24-1](#)

get_index_dataguide, DBMS_JSON PL/SQL
 function, [20-7](#), [20-9](#), [20-16](#), [20-24](#)

get() method, PL/SQL object types, [22-1](#)

H

has substring filter predicate, SQL/JSON path
 expressions, [13-2](#)

hidden virtual columns projected from JSON
 data, [20-23](#)

I

IGNORE ON MISMATCH clause, json_value,
[14-9](#)

IM column store, [27-1](#)

In-Memory Column Store, [27-1](#)
 populating JSON into, [27-3](#)
 upgrading tables with JSON data for, [27-5](#)

index, array, [13-2](#)

indexing JSON data, [26-1](#)
 composite B-tree index for multiple fields,
[26-9](#)
 for json_exists queries, [26-6](#)
 for json_table queries, [26-5](#)
 for search, [26-10](#)
 full-text and numeric-range, [26-10](#)
 function-based, [26-3](#)
 for geographic data, [24-1](#)

GeoJSON, [24-1](#)

is (not) json SQL/JSON condition, [26-2](#)

json_exists SQL/JSON condition, [26-3](#)

json_value SQL/JSON function, [26-3](#)
 data type considerations, [26-8](#)
 for geographic data, [24-1](#)
 for json_exists queries, [26-6](#)
 for json_table queries, [26-5](#)
 spatial, [24-1](#)

inserting JSON data into a column, [9-1](#)

introspection of PL/SQL object types, [22-1](#)

is json SQL/JSON condition, [5-1](#)
 and JSON null, [2-2](#)
 indexing, [26-2](#)
 STRICT keyword, [5-5](#)

is not json SQL/JSON condition, [5-1](#)
 and JSON null, [2-2](#)
 indexing, [26-2](#)
 STRICT keyword, [5-5](#)

ISO 8601 date and time formats, [13-14](#)

item method, SQL/JSON path expressions, [13-2](#),
[13-11](#)
 data type compatibility, [13-11](#)
 syntax, [B-1](#)

items data-guide field (JSON Schema keyword),
[20-9](#)

J

JavaScript array, [2-2](#)

JavaScript notation compared with JSON, [2-1](#)

JavaScript object, [2-2](#)

JavaScript object literal, [2-2](#)

JavaScript Object Notation (JSON), [2-1](#)

JSON, [2-1](#)
 character encoding, [6-1](#)
 character-set conversion, [6-1](#)
 compared with JavaScript notation, [2-1](#)
 compared with XML, [2-4](#)
 overview, [1-1](#), [2-1](#)
 support by Oracle Database, restrictions, [A-1](#)

JSON (*continued*)

- syntax, [1-1](#), [2-1](#), [2-2](#)
 - basic path expression, [13-2](#), [B-1](#)
 - strict and lax, [5-2](#)
- JSON columns, [1-2](#)
- JSON data guide, [20-1](#)
 - overview, [20-2](#)
- JSON generation functions, [21-1](#)
- JSON object types, PL/SQL
 - overview, [22-1](#)
- JSON Schema, [20-1](#)
 - keywords, [20-9](#)
- JSON search index, [26-10](#)
- json_array SQL/JSON function, [21-10](#)
- JSON_ARRAY_T PL/SQL object type, [22-1](#)
- json_arrayagg SQL/JSON function, [21-13](#)
- json_dataguide SQL function, [20-7](#), [20-9](#)
 - as an aggregate function, [20-35](#)
 - hierarchical format, [20-47](#)
 - pretty-print format, [20-47](#)
- JSON_ELEMENT_T PL/SQL object type, [22-1](#)
- json_equal Oracle SQL condition, [1](#)
- json_exists SQL/JSON condition, [15-1](#)
 - as json_table, [15-4](#)
 - indexing, [26-2](#), [26-3](#), [26-6](#)
- JSON_KEY_LIST PL/SQL object type, [22-1](#)
- json_mergepatch SQL function, [9-1](#)
- json_mergepatch() PL/SQL object-type method, [9-1](#)
- json_object SQL/JSON function, [21-6](#)
- JSON_OBJECT_T PL/SQL object type, [22-1](#)
- json_objectagg SQL/JSON function, [21-12](#)
- json_query SQL/JSON function, [17-1](#)
 - as json_table, [17-2](#)
- JSON_SCALAR_T PL/SQL object type, [22-1](#)
- json_serialize Oracle SQL function, [19-1](#)
- json_table SQL/JSON function, [18-1](#)
 - EXISTS keyword, [18-1](#)
 - FORMAT JSON keywords, [18-5](#)
 - generalizes other SQL/JSON functions and conditions, [18-8](#)
 - indexing for queries, [26-5](#)
 - NESTED PATH clause, [18-9](#)
 - PATH clause, [18-1](#)
- json_textcontains SQL/JSON condition, [26-10](#)
- json_value SQL/JSON function, [16-1](#)
 - as json_table, [16-7](#)
 - data type considerations for indexing, [26-8](#)
 - function-based indexing, [26-3](#)
 - for geographic data, [24-1](#)
 - indexing for json_exists queries, [26-6](#)
 - indexing for json_table queries, [26-5](#)
 - null JSON value, [16-4](#)
 - returning an object-type instance, [16-4](#)

K

key, JSON object
 See field, JSON object

keywords
 JSON Schema, [20-9](#)

L

lax JSON syntax, [5-2](#)

- specifying, [5-5](#)

length() item method, SQL/JSON path expressions, [13-11](#)

like filter predicate, SQL/JSON path expressions, [13-2](#)

like_regex filter predicate, SQL/JSON path expressions, [13-2](#)

limitations, Oracle Database support for JSON, [A-1](#)

loading JSON data into the database, [9-1](#)

LOB storage, using for JSON data, [3-1](#)

lower() item method, SQL/JSON path expressions, [13-11](#)

M

materialized view of JSON data, [18-12](#)

MISSING DATA clause, ON MISMATCH clause, json_value, [14-9](#)

multiple data guides for the same JSON column, [20-35](#)

N

NESTED clause, instead of json_table, [18-4](#)

NESTED PATH clause, json_table, [18-9](#)

NoSQL databases, [1-2](#)

null handling, SQL/JSON generation functions, [21-1](#), [21-3](#)

NULL ON EMPTY clause, SQL/JSON query functions, [14-8](#)

NULL ON MISMATCH clause, json_value, [14-9](#)

NULL ON NULL, SQL/JSON generation functions, [21-1](#)

NULL-handling clause, SQL/JSON generation functions, [21-1](#)

number() item method, SQL/JSON path expressions, [13-11](#)

numberOnly() item method, SQL/JSON path expressions, [13-11](#)

numeric-range indexing, [26-10](#)

O

o:frequency data-guide field, [20-9](#)

o:hidden data-guide field, [20-23](#)
o:high_value data-guide field, [20-9](#)
o:last_analyzed data-guide field, [20-9](#)
o:length data-guide field, [20-9](#)
o:low_value data-guide field, [20-9](#)
o:num_nulls data-guide field, [20-9](#)
o:path data-guide field, [20-9](#)
o:preferred_column_name data-guide field, [20-9](#)
object literal, Javascript, [2-2](#)
object member, JSON, [2-2](#)
object step, SQL/JSON path expressions, [13-2](#)
 syntax, [B-1](#)
object, Javascript and JSON, [2-2](#)
ON EMPTY clause, SQL/JSON query functions, [14-8](#)
ON MISMATCH clause, json_value, [14-9](#)
oneOf data-guide field (JSON Schema keyword), [20-9](#)
Oracle SQL conditions, [1](#)
 json_equal, [1](#)
 See also SQL/JSON conditions
Oracle SQL functions
 json_dataguide, [20-7](#), [20-9](#)
 as an aggregate function, [20-35](#)
 hierarchical format, [20-47](#)
 pretty-print format, [20-47](#)
 json_mergepatch, [9-1](#)
 json_serialize, [19-1](#)
Oracle support for JSON in the database, [A-1](#)

P

parent COLUMNS clause, json_table SQL/JSON function, [18-5](#)
parsing of JSON data to PL/SQL object types, [22-1](#)
PASSING clause, json_exists, [15-1](#)
PATH clause, json_table, [18-1](#)
path expression, SQL/JSON, [13-1](#)
 comparison, types, [13-16](#)
 item methods, [13-11](#)
 syntax, [13-2](#), [B-1](#)
performance tuning, [25-1](#)
PL/SQL functions
 DBMS_JSON.get_index_dataguide, [20-7](#), [20-9](#), [20-16](#), [20-24](#)
PL/SQL object types
 overview, [22-1](#)
PL/SQL object-type methods, [22-1](#)
 json_mergepatch(), [9-1](#)
PL/SQL procedures
 DBMS_JSON.add_virtual_columns, [20-23](#), [20-24](#), [20-27](#)
 DBMS_JSON.create_view, [20-15](#), [20-16](#)

PL/SQL procedures (*continued*)
 DBMS_JSON.create_view_on_path, [20-15](#), [20-19](#)
 DBMS_JSON.drop_virtual_columns, [20-23](#), [20-30](#)
 DBMS_JSON.rename_column, [20-9](#)
PRETTY keyword, SQL functions, [14-1](#)
pretty-printing
 in book examples, [xvi](#)
PRETTY, package DBMS_JSON, [20-16](#), [20-41](#), [20-47](#)
projecting virtual columns from JSON fields, [20-23](#)
properties data-guide field (JSON Schema keyword), [20-9](#)
property, JSON object
 See field, JSON object
put() method, PL/SQL object types, [22-1](#)

Q

queries, dot notation, [12-1](#)
 use with json_table SQL/JSON function, [18-1](#)

R

range specification, array, [13-2](#)
rawtohex SQL function, for insert or update with BLOB JSON column, [3-1](#)
relational database with JSON data, [1-2](#)
relative path expression, [13-2](#)
 syntax, [B-1](#)
rename_column, DBMS_JSON PL/SQL procedure, [20-9](#)
rendering of JSON data, [14-1](#)
restrictions, Oracle Database support for JSON, [A-1](#)
RETURNING clause
 SQL query functions, [14-1](#)
 SQL/JSON generation functions, [21-1](#)
row source, JSON
 definition, [18-1](#)

S

scalar value, JSON, [2-2](#)
scalar, SQL/JSON path expressions
 syntax, [B-1](#)
schema, JSON, [20-1](#)
schemaless database data, [1-2](#)
SDO_GEOMETRY, [24-1](#)
searching JSON data, [26-10](#)
SELECT statement, NESTED clause instead of json_table, [18-4](#)

- serialization
 - of JSON data from queries, [14-1](#)
 - of JSON data in PL/SQL object types, [22-1](#)
 - serializing JSON data, [19-1](#)
 - setting values in PL/SQL object types, [22-1](#)
 - sharding, data-guide information in index, [20-4](#)
 - sibling COLUMNS clauses, json_table SQL/JSON function, [18-5](#)
 - simple dot-notation access to JSON data, [12-1](#)
 - use with json_table SQL/JSON function, [18-1](#)
 - Simple Oracle Document Access (SODA), [1-1](#)
 - simplified syntax
 - See simple dot-notation access to JSON data
 - size() item method, SQL/JSON path expressions, [13-11](#)
 - SODA, [1-1](#)
 - spatial JSON data, [24-1](#)
 - SQL functions
 - json_dataguide, [20-7](#), [20-9](#)
 - as an aggregate function, [20-35](#)
 - hierarchical format, [20-47](#)
 - pretty-print format, [20-47](#)
 - json_mergepatch, [9-1](#)
 - SQL NESTED clause, instead of json_table, [18-4](#)
 - SQL/JSON conditions, [1](#)
 - is (not) json, [5-1](#)
 - is json
 - and JSON null, [2-2](#)
 - indexing, [26-2](#)
 - is not json
 - and JSON null, [2-2](#)
 - indexing, [26-2](#)
 - json_exists, [15-1](#)
 - as json_table, [15-4](#)
 - indexing, [26-2](#), [26-3](#)
 - json_textcontains, [26-10](#)
 - See also Oracle SQL conditions
 - SQL/JSON functions
 - for generating JSON, [21-1](#)
 - json_array, [21-10](#)
 - json_arrayagg, [21-13](#)
 - json_object, [21-6](#)
 - json_objectagg, [21-12](#)
 - json_query, [17-1](#)
 - as json_table, [17-2](#)
 - json_table, [18-1](#)
 - json_value, [16-1](#)
 - as json_table, [16-7](#)
 - function-based indexing, [24-1](#), [26-3](#)
 - null JSON value, [16-4](#)
 - returning an object-type instance, [16-4](#)
 - SQL/JSON generation functions, [21-1](#)
 - input SQL values, [21-3](#)
 - SQL/JSON path expression, [13-1](#)
 - comparison, types, [13-16](#)
 - SQL/JSON path expression (*continued*)
 - item methods, [13-11](#)
 - syntax, [13-2](#)
 - array step, [B-1](#)
 - basic, [13-2](#), [B-1](#)
 - compare predicate, [B-1](#)
 - comparison, [B-1](#)
 - condition, [B-1](#)
 - conjunction, [B-1](#)
 - disjunction, [B-1](#)
 - field name, [B-1](#)
 - filter, [B-1](#)
 - function step, [B-1](#)
 - item method, [B-1](#)
 - object step, [B-1](#)
 - relaxed, [13-9](#)
 - scalar, [B-1](#)
 - variable, [B-1](#)
 - SQL/JSON query functions
 - WITH WRAPPER keywords, [14-4](#)
 - SQL/JSON variable, [13-2](#)
 - starts with filter predicate, SQL/JSON path expressions, [13-2](#)
 - step, SQL/JSON path expressions, [13-2](#)
 - storing and managing JSON data, overview, [3-1](#)
 - strict JSON syntax, [5-2](#)
 - specifying, [5-5](#)
 - STRICT keyword
 - is (not) json SQL/JSON condition, [5-5](#)
 - SQL/JSON generation functions, [21-1](#)
 - string() item method, SQL/JSON path expressions, [13-11](#)
 - stringOnly() item method, SQL/JSON path expressions, [13-11](#)
 - support for JSON, Oracle Database, [A-1](#)
 - syntax diagrams, basic SQL/JSON path expression, [B-1](#)
- ## T
-
- tables with JSON data, [1-2](#)
 - time and date formats, ISO 8601, [13-14](#)
 - timestamp() item method, SQL/JSON path expressions, [13-11](#)
 - tree-like representation of JSON data, [22-1](#)
 - trigger for data-guide changes, [20-30](#)
 - TRUNCATE keyword, Oracle extension for SQL VARCHAR2 return value, [14-1](#)
 - TRUNCATE keyword, Oracle extension for SQL/JSON VARCHAR2 return value, [18-1](#)
 - type data-guide field (JSON Schema keyword), [20-9](#)
 - TYPE ERROR clause, ON MISMATCH clause, json_value, [14-9](#)

type() item method, SQL/JSON path expressions, [13-11](#)

U

UNCONDITIONAL keyword, SQL/JSON query functions, [14-4](#)

unique field names in JSON objects, [5-2](#)

updating JSON data, [9-1](#)

upper() item method, SQL/JSON path expressions, [13-11](#)

USER_JSON_COLUMNS view, [4-3](#)

USER_JSON_DATAGUIDE_FIELDS view, [20-12](#)

USER_JSON_DATAGUIDES view, [20-12](#)

user-defined data-guide change trigger, [20-32](#)

V

value, JSON, [2-2](#)

variable, SQL/JSON path expressions, [13-2](#)
syntax, [B-1](#)

view

create based on a data guide, [20-16](#)

create based on data guide-enabled index and a path, [20-19](#)

create using SQL/JSON function json_table, [18-12](#)

views

ALL_JSON_COLUMNS, [4-3](#)

ALL_JSON_DATAGUIDE_FIELDS, [20-12](#)

ALL_JSON_DATAGUIDES, [20-12](#)

views (*continued*)

DBA_JSON_COLUMNS, [4-3](#)

DBA_JSON_DATAGUIDE_FIELDS, [20-12](#)

DBA_JSON_DATAGUIDES, [20-12](#)

USER_JSON_COLUMNS, [4-3](#)

USER_JSON_DATAGUIDE_FIELDS, [20-12](#)

USER_JSON_DATAGUIDES, [20-12](#)

virtual columns for JSON fields, adding, [20-23](#)

based on a data guide-enabled search index, [20-27](#)

based on a hierarchical data guide, [20-24](#)

W

well formed JSON data, [5-1](#)

well-formed JSON data

ensuring, [3-1](#)

WITH UNIQUE KEYS keywords, JSON condition is json, [5-2](#)

WITH WRAPPER keywords, SQL/JSON query functions, [14-4](#)

WITHOUT UNIQUE KEYS keywords, JSON condition is json, [5-2](#)

wrapper clause, SQL/JSON query functions, [14-4](#)

WRAPPER keyword, SQL/JSON query functions, [14-4](#)

X

XML

compared with JSON, [2-4](#)

DOM, [22-1](#)