

Oracle® NoSQL Database

SQL Reference Guide



Release 26.1

F14605-33

April 2026

ORACLE®

Copyright © 2011, 2026, Oracle and/or its affiliates.

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

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

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

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

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

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

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Contents

1 Introduction to SQL for Oracle NoSQL Database

SQL Program	1
EBNF Syntax	2
Comments	3
Identifiers	3
Literals	4
Operator Precedence	5
Reserved Words	5
Case Sensitivity	5
Constraints	5

2 Oracle NoSQL Database Data Model

Atomic Data Types	1
Complex Data Types	3
JSON Data Type	5
Wildcard Data Types	6
Data Type Hierarchy	7
Data Type Definitions	9

3 Namespace Management

CREATE NAMESPACE Statement	1
SHOW NAMESPACES Statement	2
DROP NAMESPACE Statement	3
Namespace Resolution	4
Namespace Scoped Privileges	4
Granting Authorization Access to Namespaces	4

4 Region Management

CREATE REGION Statement	1
SET LOCAL REGION Statement	1
SHOW REGIONS Statement	2

5 Table Management

CREATE TABLE Statement	1
SHOW TABLES Statement	13
DESCRIBE TABLE Statement	15
Table Hierarchies	19
Example: Hierarchical Table	23
Using JSON Collection Tables	25
Using the IDENTITY Column	28
Creating Tables With an IDENTITY Column	31
Using the UUID data type	34
Using the MR_COUNTER datatype	36
Using CRDT datatype in a multi-region table	36
Create table using MR_COUNTER datatype	38
Enabling Before-images During Table Creation	40
Sequence Generator	41
DROP TABLE Statement	47
ALTER TABLE Statement	48
Altering an IDENTITY Column	53
Add or Remove a UUID column	55
Add or Remove an IDENTITY column	56
Add or Remove an MR_COUNTER column	57

6 SQL Query Management

Expressions	1
Sequences	3
Sequence Types	5
Variable Declaration	7
SELECT Expression	8
FROM Clause	8
WHERE Clause	12
GROUP BY Clause	14
Using Aggregate Functions	17
Sequence Aggregate Functions	27
ORDER BY Clause	31
SELECT Clause	36
LIMIT Clause	39
OFFSET Clause	39
Path Expressions	40

Field Step Expressions	40
Map-Filter Step Expressions	43
Array-Filter Step Expressions	44
Array-Slice Step Expressions	46
Comparison Expressions	47
Logical Operators: AND, OR, and NOT	47
IS NULL and IS NOT NULL Operators	49
Value Comparison Operators	50
Sequence Comparison Operators	53
BETWEEN Operator	56
IN Operator	59
Regular Expression Conditions	60
Exists Operator	71
Is-Of-Type Operator	72
Concatenation Operator	72
Arithmetic Expressions	73
Primary Expressions	75
Parenthesized Expressions	76
Constant Expressions	76
Column References	77
Variable References	78
Array and Map Constructors	79
Case Expressions	80
Cast Expression	82
Extract Expressions	84
Function Calls	85
Sequence Transform Expressions	86
Unnest Arrays & Maps	89
Example: Using unnesting with a GROUP BY clause	97
Joins	99
Using NESTED TABLES clause to query multiple tables in the same hierarchy	99
Example: Using NESTED TABLES clause to query multiple tables in the same hierarchy	103
Left Outer Join (LOJ)	104
Different scenarios of using an LOJ	105
Limitations of LOJ	109
Nested tables Vs LOJ	110
Example: Using Left Outer Joins	111
Inner Join	113
Different Scenarios of Using an Inner Join	115
Example: Using Inner Join	118
Inner Join vs LOJ vs NESTED TABLES	119

7 Data Row Management

INSERT Statement	1
Inserting Rows into JSON Collection Tables	5
Inserting Rows with an IDENTITY Column	8
Inserting rows into a table with a UUID column	10
Inserting rows with an MR_COUNTER column	12
Upsert statement	15
Updating rows of a table with a UUID column	19
DELETE Statement	20
UPDATE Statement	21
Update Clauses	22
SET Clause	23
ADD Clause	24
PUT Clause	25
REMOVE Clause	25
JSON MERGE Clause	25
SET TTL Clause	26
Updating rows with an IDENTITY Column	27
Updating rows with an MR_COUNTER column	27
Example: Updating Rows	30
Example: Updating JSON Data	35
Example: Updating JSON collection tables	44
Example: Updating TTL	45
Example: Updating IDENTITY defined as GENERATED ALWAYS	46
Example: Updating IDENTITY defined as GENERATED BY DEFAULT	46
JSON Collection Table Example	47

8 Indexes

About Indexes	1
CREATE INDEX Statement	2
Classification of Indexes	3
Single Field Index	4
Composite Index	4
Fixed Schema Index	5
JSON Index	5
Simple Index	6
Multikey Index	7
Nested Multikey Index	7
Composite Multikey Index	7
Specifications & Restrictions on Multikey Indexes	9

Index on JSON Collection Tables	11
Indexes on Functions	11
Examples of creating indexes on functions	12
SHOW INDEXES Statement	15
DESCRIBE INDEX Statement	16
DROP INDEX Statement	18
Appendix	19

9 Query Optimization

Using Indexes for Query Optimization	1
Finding Applicable Indexes	2
Examples: Using Indexes for Query Optimization	3
Optimizing unnesting queries with the UNNEST clause	8
Choosing the Best Applicable Index	11
Appendix	12

10 Query Plan

Overview of a query plan	1
Examples of query execution plan	3
Example 1 : Using a covering index in a query plan with only index scans	4
Example 2 : Using a covering index in a query plan with index scans and index predicates	6
Example 3: Using a non-covering index in a query plan with index scans	9
Example 4: Sort the data using a Covering index	12
Example 5: Sort the data using a field not part of the index	15
Example 6: Group the data using a Covering index	18
Example 7: Group data with fields not part of the index	22

11 GeoJson Data Management

About GeoJson Data	1
Lines and Coordinate System	4
Restrictions on GeoJson Data	6
Searching for GeoJson Data	6
Indexing GeoJson Data	12

12 Built-in Functions

Functions on Complex Values	2
Functions on Sequences	3
Functions on Timestamps	6

timestamp_add function	7
timestamp_diff function	9
get_duration function	11
timestamp_ceil function	13
timestamp_floor or timestamp_trunc function	14
timestamp_round function	16
timestamp_bucket function	17
format_timestamp function	19
parse_to_timestamp function	21
to_last_day_of_month function	23
Timestamp extract functions	24
Date extract functions	24
week/isoweek functions	26
Timestamp index extract functions	27
current_time_millis function	29
current_time function	29
Supported units	31
Function to generate a UUID string	39
Functions on Rows	41
Functions on GeoJson Data	51
Functions on Strings	53
substring Function	54
concat Function	56
upper Function	57
lower Function	58
trim Function	59
ltrim Function	61
rtrim Function	62
length Function	62
contains Function	63
starts_with Function	64
ends_with Function	65
index_of Function	66
replace Function	67
reverse Function	69
Function to Convert String to JSON	70
Functions of Mathematical Operations	73
abs function	74
arc cosine function	75
arc sine function	76
arc tan function	77
arc tan2 function	77

ceil function	78
cosine function	79
cotangent function	79
degrees function	80
Euler function	81
exp function	81
floor function	82
ln function	82
log function	83
log10 function	84
pi function	84
power function	85
radians function	85
random function	86
round function	87
sign function	88
sine function	89
square root function	89
tangent function	90
truncate function	90
Additional Examples	91

A Introduction to the SQL for Oracle NoSQL Database Shell

Running the SQL Shell	A-1
Configuring the shell	A-2
Shell Utility Commands	A-3
connect	A-4
consistency	A-5
describe	A-5
durability	A-7
exit	A-7
help	A-7
history	A-7
import	A-8
load	A-8
mode	A-9
output	A-13
page	A-13
show faults	A-13
show ddl	A-13
show indexes	A-14

show namespaces	A-14
show query	A-15
show regions	A-15
show roles	A-16
show tables	A-16
show users	A-17
timeout	A-17
timer	A-18
verbose	A-18
version	A-19

Preface

This document is intended to describe the Structured Query Language (SQL) supported by Oracle NoSQL Database. Oracle NoSQL Database can be accessed through multiple interfaces, including language SDKs and tools. SQL provides a common, declarative way to define, query, and manipulate data stored in Oracle NoSQL Database. Understanding the language is essential for designing schemas, writing efficient queries, and building reliable data access layers.

This guide begins with an introduction to SQL as implemented in Oracle NoSQL Database. It covers SQL program structure, notation conventions (including EBNF syntax), and key syntax rules for identifiers, literals, comments, operator precedence, reserved words, case sensitivity, and constraints. These topics provide a baseline for reading and writing SQL statements correctly.

The guide then explains the Oracle NoSQL Database data model, covering atomic types, complex types, JSON types, and wildcard types. Together, these types enable modeling and manipulating both structured and hierarchical data.

Building on this conceptual foundation, the remainder of the guide provides detailed syntax, semantics, and examples for the SQL operations supported by Oracle NoSQL Database.

Audience

This book is intended for developers, database designers, and users who want to write SQL statements for Oracle NoSQL Database. It serves as both a learning resource for new users and a reference manual for experienced developers.

For information about using SQL with the language SDKs to build applications, see – *Developers Guide*.

1

Introduction to SQL for Oracle NoSQL Database

Structured Query Language (SQL) is the set of statements with which all programs and users access data in the Oracle NoSQL Database. This book provides information on SQL as used by Oracle NoSQL Database. Application programs and Oracle tools often allow users access to the Oracle NoSQL Database without using SQL directly.

This chapter contains the following topics:

- [SQL Program](#)
- [EBNF Syntax](#)
- [Comments](#)
- [Identifiers](#)
- [Literals](#)
- [Operator Precedence](#)
- [Reserved Words](#)
- [Case Sensitivity](#)
- [Constraints](#)

Note

No prior knowledge of SQL is required for reading this document.

SQL Program

The data model of Oracle NoSQL Database supports (a) flat relational data, (b) hierarchical typed (schema-full) data, and (c) schema-less JSON data. SQL for Oracle NoSQL Database is designed to handle all such data in a seamless fashion, without any *impedance mismatch* among the different sub models.

In the current version, an SQL program consists of a single statement, which can be a non-updating query (read-only DML statement), an updating query (updating DML statement), a data definition command (DDL statement), a user management and security statement, or an informational statement. This is illustrated in the following syntax, which lists all the statements supported by the current SQL version.

```
program ::=  
(  
    query |  
    insert\_statement |  
    Upsert statement |  
    delete\_statement |  
    update\_statement |
```

```

create\_namespace\_statement |
show\_namespaces\_statement |
drop\_namespace\_statement |
create\_region\_statement |
show\_regions\_statement |
drop\_region\_statement |
create\_table\_statement |
show\_tables\_statement |
describe\_table\_statement |
alter\_table\_statement |
drop\_table\_statement |
create\_index\_statement |
show\_indexes\_statement |
describe\_index\_statement |
drop\_index\_statement |
create_text_index_statement |
create_user_statement |
create_role_statement |
drop_role_statement |
drop_user_statement |
alter_user_statement |
grant_statement |
revoke_statement |
) EOF

```

This document is concerned with the first 19 statements in the above list, that is, with read-only queries, insert/delete/update statements, namespace statements, and DDL statements, excluding text indexes. The document describes the syntax and semantics for each statement, and supplies examples. The programmatic APIs available to compile and execute SQL statements and process their results are described in the *Java Direct Driver Developer's Guide*.

EBNF Syntax

This specification uses EBNF meta-syntax to specify the grammar of SQL for Oracle NoSQL Database. The following EBNF notations apply:

- **Upper-case words** are used to represent keywords, punctuation characters, operator symbols, and other syntactical entities that are recognized by EBNF as terminals (aka tokens) in the query text. For example, SELECT stands for the "select" keyword in the query text. Notice that keywords are case-insensitive; "select" and "sELECT" are both the same keyword, represented by the SELECT terminal.
- **Lower-case words** are used for non-terminals. For example, array_step : [expression] means that array_step is an expression enclosed in square brackets.
- " " Anything enclosed in quotes is also considered a terminal. For example, the following production rule defines the value-comparison operators as one of the =, >= symbols: For example, val_comp : "=" | ">=" .
- * means 0 or more of whatever precedes it. For example, field_name* means 0 or more field names.
- + means 1 or more of whatever precedes it. For example, field_name+ means 1 or more field names.

- `[]` means optional, i.e., zero or 1 of whatever contained in it. For example, `[field_name]` means zero or one field names.
- `|` means this or that. For example, `INT | STRING` means an integer, or a string.
- `()` Parentheses are used to group EBNF sub-expressions together. For example, `(INT | STRING) [comment]` means an integer, or a string, followed by a comment, or just an integer, or a string, followed by nothing.

Comments

The language supports comments in both DML and DDL statements. Such comments have the same semantics as comments in a regular programming language, that is, they are not stored anywhere, and have no effect to the execution of the statements. The following comment constructs are recognized:

`/* comment */`

Potentially multi line comment.

Note

However, if a '+' character appears immediately after the opening `/*`, and the comment is next to a `SELECT` keyword, the comment is actually not a comment but a hint for the query processor. See [Choosing the Best Applicable Index](#).

`// comment`

Single line comment.

`# comment`

Single line comment.

As we will see, DDL statements may also contain comment clauses, which are stored persistently as properties of the created data entities. Comment clauses start with the `COMMENT` keyword, followed by a string literal, which is the content of the comment.

Syntax

```
comment ::= COMMENT string
```

Identifiers

An identifier is a sequence of characters conforming to the following rules:

- It starts with a latin alphabet character (characters 'a' to 'z' and 'A' to 'Z').
- The characters after the first one may be any combination of latin alphabet characters, decimal digits ('0' to '9'), or the underscore character ('_').
- It is not one of the reserved words. The only reserved words are the literals `TRUE`, `FALSE`, and `NULL`.

ID is the terminal that represents the identifiers. However, in the grammar rules presented in this document we will use the non-terminal symbol `id` to denote identifiers.

Syntax

```

ID ::= ALPHABET (ALPHABET | DIGIT | '_' ) *
ALPHABET ::= 'a'..'z'|'A'..'Z'
DIGIT ::= '0'..'9'

id ::=
ID |
ADD | ALTER | ANCESTORS | AND |
ANY | ANYATOMIC | ANYJSONATOMIC | ANYRECORD |
ARRAY | AS | ASC |
BINARY | BOOLEAN | BY |
CASCADE | CASE | CAST | COMMENT | COUNT | CREATE |
DAYS | DECLARE | DEFAULT | DESC | DESCENDANTS | DISTINCT | DOUBLE | DROP |
ELSE | END | ENUM | EXISTS | EXTRACT | FIRST | FLOAT | FROM |
GEOMETRY | GROUP | HOURS | IF | IN | INDEX | INTEGER | IS |
JSON | KEY | KEYS |
LAST | LIMIT | LONG | MAP |
NAMESPACE | NESTED | NO | NOT | NULLS |
OF | OFFSET | ON | OR | ORDER |
POINT | PRIMARY | RECORD |
SELECT | SHARD | STRING |
TABLE | TABLES | THEN | TTL | TYPE |
USING | VALUES |
WHEN | WHERE | WITH

```

Literals

A literal (a.k.a constant value) is a fixed value appearing in the query text. There are four kinds of literals: numbers, strings, boolean values, and the JSON NULL value. The following production rules are used to recognize literals in the query text. The [Constant Expressions](#) section describes how the tokens listed below are translated into instances of the data model.

Syntax

```

INT_CONSTANT ::= DIGIT+
FLOAT_CONSTANT ::=
    (DIGIT* '.' DIGIT+ [(E|e) [+|-] DIGIT+]) |
    (DIGIT+ (E|e) [+|-] DIGIT+)
NUMBER_CONSTANT ::= (FLOAT_CONSTANT | INT_CONSTANT) (N|n)
STRING_CONSTANT ::= '\'' [(ESC|.)*] '\''
DSTRING_CONSTANT ::= '"' [(ESC|.)*] '"'
ESC ::= '\\' ([\\\/bfnrt]|UNICODE)
DSTR_ESC ::= '"' ([\\\/bfnrt]|UNICODE)
UNICODE ::= 'u' HEX HEX HEX HEX
TRUE ::= (T|t)(R|r)(U|u)(E|e)
FALSE ::= (F|f)(A|a)(L|l)(S|s)(E|e)
NULL ::= (N|n)(U|u)(L|l)(L|l)

```

Note

The literals TRUE, FALSE, and NULL are reserved words.

Operator Precedence

The relative precedence among the various operators and expressions in SQL for Oracle NoSQL Database is defined implicitly by the order in which the grammar rules for these operators and expressions are listed in the grammar specification. Specifically, the earlier a grammar rule appears, the lower its precedence. For example, consider the following 3 rules that define the syntax for the OR, AND, and NOT operators. Because `or_expr` appears before `and_expr` and `not_expr`, OR has lower precedence than AND and NOT. And AND has lower precedence than NOT, because `and_expr` appears before `not_expr`. As a result, an expression like `a = 10 and not b > 5 or c < 20 and c > 1` is equivalent to `(a = 10 and (not b > 5)) or (c < 20 and c > 1)`. See [Logical Operators: AND, OR, and NOT](#) for more details.

```
or_expression ::= and_expression | (or_expression OR and_expression)
and_expression ::= not_expression | (and_expression AND not_expression)
not_expression ::= [NOT] is_null_expression
```

Reserved Words

Reserved words are words that look like identifiers, but cannot be used as identifiers (i.e., in places where identifiers are expected). SQL for Oracle NoSQL Database has a short list of reserved words. Currently, this list consists of the following (case-insensitive) words: **TRUE**, **FALSE**, and **NULL**.

Case Sensitivity

In Oracle NoSQL Database, the following fields are not case-sensitive while using in SQL commands.

- table name
- namespace name
- field name
- secondary-index name
- text-index name
- region-name

Constraints

There is a limitation on the length of the name of the table or index that you want to create.

Table 1-1 Constraints on the length and content

Description	Constraints on the length	Constraints on the content
table name	256	Must contain only alphanumeric characters or "_" and must start with an alphabetic character
namespace name	128	Must contain only alphanumeric characters or "_" and must start with an alphabetic character

Table 1-1 (Cont.) Constraints on the length and content

Description	Constraints on the length	Constraints on the content
field name	64	No constraints
secondary-index name	64	Must contain only alphanumeric characters or "_" and must start with an alphabetic character
text-index name	64	Must contain only alphanumeric characters or "_" and must start with an alphabetic character
region-name	128	Must contain only alphanumeric characters or "_" and must start with an alphabetic character

Oracle NoSQL Database Data Model

In Oracle NoSQL Database, data is modeled as typed items. A **typed item** (or simply **item**) is a pair consisting of a data type and a value. A **type** is a set of possible values that may be stored in a database field (e.g. a table column) or be returned by some computation (e.g. during the evaluation of a query). In any item, the item value must be an instance of the item type. The values in this set are called the instances of the type. An item is said to be an instance of a type T if its value is an instance of T.

The Oracle NoSQL Database data model consists of various data types that allow for the storage and manipulation of hierarchical data. The Oracle NoSQL Database data types can be broadly classified into **atomic types**, **complex types**, **JSON types**, and **wildcard types**. Instances of atomic data types are single, indivisible units of data. Instances of complex data types contain multiple items and provide access to their nested items. JSON types contain the JSON objects in the name/value format. Wildcard types are similar to abstract supertypes in object-oriented programming languages. They combine instances of other data types and they are used to support data that do not have a fixed structure. This chapter describes each of these data types in detail.

Note

Data types describe the kind of data that can be stored in an Oracle NoSQL Database, as well as the kind of data generated during the evaluation of a query.

This chapter contains the following topics:

- [Atomic Data Types](#)
- [Complex Data Types](#)
- [JSON Data Type](#)
- [Wildcard Data Types](#)
- [Data Type Hierarchy](#)
- [Data Type Definitions](#)

Atomic Data Types

An instance of an atomic data type is a single, indivisible unit of data. The following table lists the atomic types currently available. For each type, a description of its instances is given.

Table 2-1 Atomic Data Types

Data Type	Description	Example
INTEGER	An integer between -2^{31} to $2^{31}-1$.	2147483647
LONG	An integer between -2^{63} to $2^{63}-1$.	9223372036854775807
FLOAT	A single precision IEEE 754 floating point number.	100.12345

Table 2-1 (Cont.) Atomic Data Types

Data Type	Description	Example
DOUBLE	A double precision IEEE 754 floating point number.	100.12345678901234
NUMBER	An arbitrary-precision signed decimal number (equivalent to the Java BigDecimal type).	100.123456789
STRING	A sequence of zero or more unicode characters.	"Oracle"
BOOLEAN	Has only two possible values. TRUE and FALSE.	TRUE
BINARY	An uninterpreted sequence of zero or more bytes.	Type: BINARY Type Instance: "0x34 0xF5"
FIXED BINARY (S)	An uninterpreted sequence of S bytes.	Type: BINARY(3) Type Instance: "0x34 0xF5 0xAB"
ENUM (T1, T2, ..., Tn)	One of the symbolic tokens (T1, T2, ..., Tn) explicitly listed in the ENUM type. The order in which the tokens are listed is important. For example, ENUM(a, b) and ENUM(b, a) are two distinct types.	Type: ENUM(Chennai, Bangalore, Boston) Type Instance: Boston
TIMESTAMP (P)	A value representing a point in time as a date (year, month, day), time (hour, minute, second), and number of fractions of a second. The scale at which fractional seconds are counted is called precision P of a timestamp. The minimum precision is 0 and maximum is 9. For example, a precision of 0 means that no fractional seconds are stored, 3 means that the timestamp stores milliseconds, and 9 means a precision of nanoseconds. There is no timezone information stored in timestamp; they are all assumed to be in the UTC timezone.	Type: timestamp(3) Type Instance : '2020-01-20T12:15:054'

In addition to the kind of atomic values described above, the Oracle NoSQL Database data model includes the following 2 atomic values:

Table 2-2 Atomic Values

Data Type	Description
JSON NULL	This is considered to be an instance of the JSON data type. For more information about JSON, see Wildcard Data Types .
SQL NULL	This is a special value that is used to indicate the fact that a value is unknown or inapplicable. NULL is assumed to be an instance of every data type.

Note

Although an instance of a numeric type may be semantically equivalent to an instance of another numeric type, the 2 instances are distinct. For example, there is a single number 3 in the universe, but there are 5 different instances of 3 in the data model, one for each of the numeric types.

Complex Data Types

An instance of a complex data type contains multiple values and provides access to its nested values. Currently, Oracle NoSQL Database supports the following kinds of complex values:

Table 2-3 Complex Data Types

Data Type	Description	Example
ARRAY (T)	In general, an array is an ordered collection of zero or more items. The items of an array are called elements. Arrays cannot contain any NULL values. An instance of ARRAY (T) is an array whose elements are all instances of type T. T is called element type of the array.	Type: ARRAY (INTEGER) Type Instance: [600004,560076,01803]
MAP (T)	In general, a map is an unordered collection of zero or more key-item pairs, where all keys are strings. The keys in a map must be unique. The key-item pairs are called fields. The keys are called fields names, and the associated items are called field values. Maps cannot contain any NULL field value. An instance of MAP (T) is a map whose field values are all instance of type T. T is called the value type of the map.	Type: MAP(INTEGER) Type Instance: { "Chennai":600004, "Bangalore":560076, "Boston":01803 }
RECORD (k1 T1 n1, k2 T2 n2, ..., kn Tn nn)	In general, a record is an ordered collection of one or more key-item pairs, where all keys are strings. The keys in a record must be unique. The key-item pairs are called fields. The keys are called fields names, and the associated items are called field values. Records may contain NULL as field value. An instance of RECORD (k1 T1 n1, k2 T2 n2, ..., kn Tn nn) is a record of exactly n fields, where for each field i (a) the field name is ki, (b) the field value is an instance of type Ti, and (c) the field conforms to the nullability property ni, which specifies whether the field value may be NULL or not. Contrary to maps and arrays, it is not possible to add or remove fields from a record. This is because the number of fields and their field names are part of the record type definition associated with a record value.	Type: RECORD(country STRING, zipcode INTEGER, state STRING, street STRING) Type Instance: { "country":"US", "zipcode":600004, "state":"Arizona", "street":"4th Block" }

Example 2-1 Complex Data Type

The following examples illustrate the difference between the way data get stored in various complex data types.

To store the zip codes of multiple cities when the number of zip codes is not known in advance, you can use arrays.

Declaration:
ARRAY(INTEGER)

Example:
[600004, 560076, 01803]

To store the names of multiple cities along with their zip codes and the number of zip codes are not known, you can use maps.

Declaration:
MAP(INTEGER)

Example:
{
 "Chennai":600004,
 "Bangalore":560076,
 "Boston":01803
}

Records are used for an ordered collection. If you want to store a zip code as part of a bigger data set, you can use records. In this example, only a zip code is stored in a record.

Declaration:
RECORD(zipcode INTEGER)

Example:
{
 "zipcode":600004
}

You can combine multiple complex data types so that more complex data can be stored. For the same zipcode data, the following example combines two complex data types and stores the information.

Declaration:
ARRAY(RECORD(area STRING, zipcode INTEGER))

Example:
[
 {"area":"Chennai","zipcode":600004},
 {"area":"Bangalore","zipcode":560076},
 {"area":"Boston","zipcode":01803}
]

Example 2-2 Complex Data Type

This example illustrate the differences in the way a map and a record should be declared for storing complex data.

Let us consider the following data.

```
{
  "name":"oracle",
  "city":"Redwood City",
  "zipcode":94065,
  "offices":["Chennai","Bangalore","Boston"]
}
```

For the above data, you declare a map and a record as shown below.

```
Record
(
  name STRING,
  city STRING,
  zipcode INTEGER,
  offices Array(STRING)
)
```

```
Map(ANY)
```

JSON Data Type

JSON is a lightweight text-based open standard designed for human-readable data interchange. It is easy to read and write, and language independent. A JSON text is data represented in name/value format. A valid JSON text contains an object surrounded by curly brackets and includes a comma-separated list of name/value pairs. Each name is followed by a ':'(colon) character. JSON data is written to JSON data columns by providing a JSON object.

Table 2-4 JSON Data Type

Data Type	Description	Examples
JSON	The JSON type represents all valid JSON values (numbers, strings, array(JSON), map(JSON), and JSON null). Specifically, an instance of JSON can be	{ "city" : "Santa Cruz", "zip" : 95008, "phones" : [{ "area" : 408, "number" : 4538955, "kind" : "work" }, { "area" : 831, "number" : 7533341, "kind" : "home" }] }
	1. an instance of ANYJSONATOMIC,	"Santa Cruz"
	2. or an array whose elements are all instances of JSON,	95008
	3. or a map whose field values are all instances of JSON.	true
		[12, "foo", { "city":"Santa Cruz"}, [2, 3]]

To load JSON data into a table, Oracle NoSQL Database offers programmatic APIs to ingest JSON text. The supported NoSQL SDKs handle this differently. Oracle NoSQL Database will

parse the input JSON text internally and map its constituent pieces to the types described as follows:

- Numbers are converted to integer, long, or double items, depending on the actual value of the number (float items are not used for JSON).
- Strings in the input text are mapped to string items.
- Boolean values are mapped to boolean items.
- The null values are mapped to JSON nulls depending on the way the input is supplied. For example, If you supply input text as "myvalue" : null, this is mapped as a JSON null. Whereas, supplying "myvalue" : "null" sets the string value to the text "null".
- When an array is encountered in the input text, an array item is created whose type is Array(JSON). This is done unconditionally, no matter what the actual contents of the array might be.
- When a JSON object is encountered in the input text, a map item is created whose type is Map(JSON), unconditionally.

In general, the result of this parsing is a tree of maps, arrays, and atomic values. For persistent storage, the tree is serialized into a binary format.

Note

There is no JSON equivalent to the TIMESTAMP data type, so if input text contains a string in the TIMESTAMP format, it is stored as a string item in the JSON column.

JSON data is schema-less, in the sense that a field of type JSON can have very different kinds of values in different table rows. For example, if info is a top-level table column of type JSON, in one row the value of info may be an integer, in another row an array containing a mix of doubles and strings, and in a third row a map containing a mix of other maps, arrays, and atomic values. Furthermore, the data stored in a JSON column or field can be updated in any way that still produces a valid JSON instance. As a result, each JSON tree (either in main memory or as a serialized byte array on disk) is self-describing about its contents.

Wildcard Data Types

The Oracle NoSQL Database data model includes the following wildcard data types:

Table 2-5 Wildcard Data Types

Data Type	Description	Examples
ANY	Any instance of any NoSQL type is an instance of the ANY type as well.	<pre>{ "city" : "Santa Cruz", "zip" : 95008, "phones" : [{ "area" : 408, "number" : 4538955, "kind" : "work" }, { "area" : 831, "number" : 7533341, "kind" : "home" }] }</pre> "Santa Cruz" 95008 TRUE '0x34 0xF5' '2020-01-20T12:15:054' [12, "foo", { "city": "Santa Cruz"}, [2, 3]]
ANYATOMIC	Any instance of any other atomic type is an instance of the ANYATOMIC type as well. The json null value is also an instance of ANYATOMIC.	"Santa Cruz" 95008 TRUE '0x34 0xF5' '2020-01-20T12:15:054'
ANYJSONATOMIC	Any instance of a numeric type, the string type, and the boolean type is an instance of the ANYJSONATOMIC type as well. The json null value is also an instance of ANYJSONATOMIC.	"Santa Cruz" 95008 true
JSON	See the <i>Description</i> column in the Table 2-4 .	See the <i>Examples</i> column in the Table 2-4 .
ANYRECORD	Any instance of any other RECORD type is an instance of the ANYRECORD type as well.	<pre>{ "city" : "Santa Cruz", "zip" : 95008 }</pre>

A data type is called **precise** if it is not one of the wildcard types and, in the case of complex types, all of its constituent types are also precise. Items that have precise types are said to be **strongly typed**.

Wildcard types are abstract, which means that no item can have a wildcard type as its type. However, items may have an imprecise type. For example, an item may have MAP(JSON) as its type, indicating that its value is a map that can store field values of different types, as long as all of these values belong to the JSON type. In fact, MAP(JSON) is the type that represents all JSON objects (JSON documents), and ARRAY(JSON) is the type that represents all JSON arrays.

Data Type Hierarchy

The Oracle NoSQL Database data model also defines a **subtype-supertype relationship** among the types presented above. The relationship can be expressed as an `is_subtype(T, S)` function that returns true if type T is a subtype of type S and false otherwise. `is_subtype(T, S)` returns true in the following cases:

- T and S are the same type. So, every type is a subtype of itself. We say that a type T is a **proper subtype** of another type S if T is a subtype of S and T is not equal to S.

- S is the ANY type. So, every type is a subtype of ANY.
- S is the ANYATOMIC type and T is an atomic type.
- S is ANYJSONATOMIC and T is one of the numeric types or the STRING type, or the BOOLEAN type.
- S is NUMBER and T is one of the other numeric types.
- S is LONG and T is INTEGER.
- S is DOUBLE and T is FLOAT.
- S is STRING and T is UUID.
- S is TIMESTAMP(p2), T is TIMESTAMP(p1) and $p1 \leq p2$.
- S is BINARY and T is FIXED_BINARY.
- S is ARRAY(T2), T is ARRAY(T1) and T1 is a subtype of T2.
- S is MAP(T2), T is MAP(T1) and T1 is a subtype of T2.
- S and T are both record types and (a) both types contain the same field names and in the same order, (b) for each field, its type in T is a subtype of its type in S, and (c) if the field is nullable in T, it is also nullable in S.
- S is JSON and T is (a) an array whose element type is a subtype of JSON, or (b) a map whose value type is a subtype of JSON, or (c) ANYJSONATOMIC or any of its subtypes.

Note

The `is_subtype` relationship is transitive, that is, if type A is a subtype of type B and B is a subtype of C, then A is a subtype of C.

The `is_subtype` relationship is important because the usual subtype-substitution rule is supported by SQL for Oracle NoSQL Database: if an operation expects input items of type T then it can also operate on items of type S, where S is a subtype of T. However, there are two exceptions to this rule:

1. DOUBLE and FLOAT are subtypes of NUMBER. However, DOUBLE and FLOAT include three special values in their domain:

- a. NaN (not a number)
- b. Positive infinity
- c. Negative infinity

These three values are not in the domain of NUMBER. You can provide DOUBLE/FLOAT types to NUMBER type as long as these are not one of the three special values; otherwise, an error will be raised.

2. Items whose type is a proper subtype of ARRAY (JSON) or MAP (JSON) cannot be used as:


- a. RECORD/MAP field values if the field type is JSON, ARRAY (JSON) or MAP (JSON)
- b. Elements of ARRAY whose element type is JSON, ARRAY (JSON) or MAP (JSON)

This is in order to disallow strongly typed data to be inserted into JSON data.

For example, consider a JSON document M, i.e., a MAP value whose associated type is a MAP (JSON). M may contain an ARRAY value A that contains only INTEGERS. However, the type associated with A cannot be ARRAY (INTEGER), it must be ARRAY (JSON). If A

had type `ARRAY (INTEGER)`, the user would not be able to add any non-`INTEGER` values to `A`, i.e., the user would not be able to update the JSON document in a way that would still keep it a JSON document.

Figure 2-1 SQL Type Hierarchy



Data Type Definitions

The Oracle NoSQL Database data model types inside SQL statements are referred to using `type_definition` syntax. This syntax is used both in data definition language (DDL) statements and data manipulation language (DML) statements.

Syntax

```

type_definition ::=
    INTEGER |
    LONG |
    FLOAT |
    DOUBLE |
    NUMBER |
  
```

```

STRING |
BOOLEAN |
ANY |
JSON |
ANYRECORD |
ANYATOMIC |
ANYJSONATOMIC |
array_definition |
map_definition |
binary_definition |
timestamp_definition |
enum_definition |
record_definition

array_definition ::= ARRAY "(" type_definition ")"
map_definition ::= MAP "(" type_definition ")"
binary_definition ::= BINARY ["(" INT_CONSTANT ")"]
timestamp_definition ::= TIMESTAMP ["(" INT_CONSTANT ")"]

enum_definition ::= ENUM "(" id_list ")"
id_list ::= id ["," id]

record_definition ::= RECORD "(" field_definition ("," field_definition)* ")"
field_definition ::= id type_definition [default_definition] [comment]
default_definition ::=
    (default_value [NOT NULL]) | (NOT NULL default_value)
default_value ::= DEFAULT (number | string | TRUE | FALSE | id)

```

Semantics

type_definition

When the `type_def` grammar rule is used in any DDL statement, the only wildcard type that is allowed is the JSON type. So, for example, it is possible to create a table with a column whose type is JSON, but not a column whose type is ANY.

timestamp_definition

The precision is optional while specifying a `TIMESTAMP` type. If omitted, the default precision is 9 (nanoseconds). This implies that the type `TIMESTAMP` (with no precision specified) is a supertype of all other `TIMESTAMP` types (with a specified precision). However, in the context of a `CREATE TABLE` statement, a precision must be explicitly specified. This restriction is to prevent users from inadvertently creating `TIMESTAMP` values with precision 9 (which takes more space) when in reality they don't need that high precision.

record_definition

Field default values and descriptions do not affect the value of a `RECORD` type, i.e., two `RECORD` types created according to the above syntax and differing only in their default values and/or field descriptions have the same value (they are essentially the same type).

field_definition

The `field_definition` rule defines a field of a `RECORD` type. It specifies the field name, its type, and optionally, a default value and a comment. The comment, if present, is stored persistently as the field's description.

default_definition

By default, all RECORD fields are nullable. The `default_definition` rule can be used to declare a field not-nullable or to specify a default value for the field. When a record is created, if no value is assigned to a field, the default value is assigned by Oracle NoSQL Database, if a default value has been declared for that field. If not, the field must be nullable, in which case the null value is assigned. Currently, default values are supported only for numeric types, STRING, BOOLEAN, and ENUM.

3

Namespace Management

A namespace defines a group of tables, within which all of the table names must be uniquely identified. This chapter describes namespaces and how to create and manage the namespaces in Oracle NoSQL Database.

Namespaces permit you to do table privilege management as a group operation. You can grant authorization permissions to a namespace to determine who can access both the namespace and the tables within it.

Namespaces permit tables with the same name to exist in your database store. To access such tables, you can use a fully qualified table name. A fully qualified table name is a table name preceded by its namespaces, followed with a colon (:), such as `ns1:table1`.

All tables are part of some namespace. There is a default Oracle NoSQL Database namespace, called `sysdefault`. All tables are assigned to the default `sysdefault` namespace, until or unless you create other namespaces, and create new tables within them. You cannot change an existing table's namespace. Tables in `sysdefault` namespace do not require a fully qualified name and can work with just the table name. For example, to access a table in `sysdefault` namespace, you can just specify `table1` instead of `sysdefault:table1`.

Note

In a store that was created new or was upgraded from a version prior to 18.3, all the tables will be part of `sysdefault` namespace.

This chapter contains the following topics:

- [CREATE NAMESPACE Statement](#)
- [SHOW NAMESPACES Statement](#)
- [DROP NAMESPACE Statement](#)
- [Namespace Resolution](#)
- [Namespace Scoped Privileges](#)
- [Granting Authorization Access to Namespaces](#)

CREATE NAMESPACE Statement

You can add a new namespace by using the `CREATE NAMESPACE` statement.

Syntax

```
create_namespace_statement ::=
    CREATE NAMESPACE [IF NOT EXISTS] namespace_name

namespace_name ::= name_path
```

```
name_path ::= field_name ( "." field_name ) *  
field_name ::= id | DSTRING
```

Semantics

IF NOT EXISTS: This is an optional clause. If you specify this clause, and if a namespace with the same name exists, then this is a noop and no error is generated. If you don't specify this clause, and if a namespace with the same name exists, an error is generated indicating that the namespace already exists.

Note

Namespace names starting with sys are reserved. You cannot use the prefix sys for any namespaces.

Example 3-1 Create Namespace Statement

The following statement defines a namespace named *ns1*.

```
CREATE NAMESPACE IF NOT EXISTS ns1
```

Create Table in the namespace ns1 as follows:

```
CREATE TABLE ns1:t (id INTEGER, name STRING, primary key (id))
```

```
INSERT INTO ns1:t VALUES (1, 'Smith')
```

```
SELECT * FROM ns1:t
```

Output:

```
{"id":1,"name":"Smith"}
```

SHOW NAMESPACES Statement

Syntax

```
show_namespaces_statement ::= SHOW [AS JSON] NAMESPACES
```

Semantics

The show namespaces statement provides the list of namespaces in the system.

AS JSON can be specified if you want the output to be in JSON format.

Example 3-2 Show Namespaces

The following statement lists the namespaces present in the system.

```
SHOW NAMESPACES
```

Output:

```
namespaces
ns1
sysdefault
```

Example 3-3 Show Namespaces

The following statement lists the namespaces present in the system in JSON format.

```
SHOW AS JSON NAMESPACES
```

Output:

```
{"namespaces" : ["ns1", "sysdefault"]}
```

DROP NAMESPACE Statement

You can remove a namespace by using the DROP NAMESPACE statement.

Syntax

```
drop_namespace_statement ::=
    DROP NAMESPACE [IF EXISTS] namespace\_name [CASCADE]
```

Semantics

IF EXISTS: This is an optional clause. If you specify this clause, and if a namespace with the same name does not exist, no error is generated. If you don't specify this clause, and if a namespace with the same name does not exist, an error is generated indicating that the namespace does not exist.

CASCADE: This is an optional clause that enables you to specify whether to drop the tables and their indexes in this namespace. If you specify this clause, and if the namespace contains any tables, then the namespace together with all the tables in this namespace will be deleted. If you don't specify this clause, and if the namespace contains any tables, then an error is generated indicating that the namespace is not empty.

Note

You cannot drop the default namespace, *sysdefault*.

Example 3-4 Drop Namespace Statement

The following statement removes the namespace named *ns1*.

```
DROP NAMESPACE IF EXISTS ns1 CASCADE
```

Namespace Resolution

To resolve a table from a *table_name* that appears in an SQL statement, the following rules apply:

- if the *table_name* contains a namespace name, no resolution is needed, because a qualified table name uniquely identifies a table.
- if you don't specify a namespace name explicitly, the namespace used is the one contained in the *ExecuteOptions* instance that is given as input to the *executeSync()*, *execute()*, or *prepare()* methods of *TableAPI*. See *Java Direct Driver Developer's Guide*.
- if *ExecuteOptions* doesn't specify a namespace, the default *sysdefault* namespace is used.

Using different namespaces in *ExecuteOptions* allows executing the same queries on separate but similar tables.

Namespace Scoped Privileges

You can add one or more namespaces to your store, create tables within them, and grant permission for users to access namespaces and tables. For general information on managing Roles and Users, see *Grant Roles or Privileges* in the *Security Guide*.

For information on implication relationship among Oracle NoSQL Database privileges, see *Privilege Hierarchy* in the *Security Guide*.

Granting Authorization Access to Namespaces

You can manage permission for users or roles to access namespaces and tables. These are the applicable permissions given to the developers and other users:

Table 3-1 Namespace Privileges and Permissions

Privilege	Description
CREATE_ANY_NAMESPACE DROP_ANY_NAMESPACE	Grant permission to a user or to a role to create or drop any namespace. GRANT CREATE_ANY_NAMESPACE TO <User Role>; GRANT DROP_ANY_NAMESPACE TO <User Role>;

Table 3-1 (Cont.) Namespace Privileges and Permissions

Privilege	Description
CREATE_TABLE_IN_NAMESPACE DROP_TABLE_IN_NAMESPACE EVOLVE_TABLE_IN_NAMESPACE	Grant permission to a user or to a role to create, drop or evolve tables in a specific namespace. You can evolve tables to update table definitions, add or remove fields, or change field properties, such as a default value. You may even add a particular kind of column, like an IDENTITY column, to increment some value automatically. Only tables that already exist in the store are candidates for table evolution. For more details, see Alter Table. <pre>GRANT CREATE_TABLE_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>; GRANT DROP_TABLE_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>; GRANT EVOLVE_TABLE_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>user_role;</pre>
CREATE_INDEX_IN_NAMESPACE DROP_INDEX_IN_NAMESPACE	Grant permission to a user or to a role to create or drop an index in a specific namespace. <pre>GRANT CREATE_INDEX_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>; GRANT DROP_INDEX_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>;</pre>
READ_IN_NAMESPACE INSERT_IN_NAMESPACE DELETE_IN_NAMESPACE	Grant permission to a role to read, insert, or delete items in a specific namespace. <pre>GRANT READ_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>; GRANT INSERT_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>; GRANT DELETE_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>;</pre>
MODIFY_IN_NAMESPACE	Helper label for granting or revoking permissions to all DDL privileges for a specific namespace to a user or role. <pre>GRANT MODIFY_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>; REVOKE MODIFY_IN_NAMESPACE ON NAMESPACE namespace_name TO <User Role>;</pre>

Grant privileges on a namespace

You can grant permissions to a role or a user on a namespace. Following is the syntax for granting permissions on a namespace:

```
GRANT {Namespace-scoped privileges} ON NAMESPACE namespace_name TO <User|Role>
Namespace-scoped privileges ::= namespace_privilege [, namespace_privilege]
```

where,

- namespace_privilege

The namespace privilege that can be granted to a user or a role. For more information on the applicable privileges, see the *Privilege* column in the [Namespace Privileges and Permissions](#) table.

- namespace_name

The namespace that the user wishes to access.

- <User|Role>

The name of the KVStore user or the role of a user.

For example, you can grant read access to a user for all the tables in the namespace.

```
GRANT READ_IN_NAMESPACE ON NAMESPACE ns1 TO Kate
```

Here, ns1 is the namespace and Kate is the user.

Note

The label MODIFY_IN_NAMESPACE can be used as a helper for granting or revoking permissions to all DDL privileges for a specific namespace to a user or role.

Revoke privileges on a namespace

You can revoke the permissions from a role or a user on a namespace. Following is the syntax for revoking the permissions on a namespace.

```
REVOKE {Namespace-scoped privileges} ON NAMESPACE namespace_name FROM <User|  
Role>  
Namespace-scoped privileges ::= namespace_privilege [, namespace_privilege]
```

where,

- namespace_privilege

The namespace privilege that can be revoked from a user or a role. For more information on the applicable privileges, see the *Privilege* column in the [Namespace Privileges and Permissions](#) table.

- namespace_name

The namespace that the user wishes to access.

- <User|Role>

The name of the KVStore user or the role of a user.

For example, you can revoke the read access from a user for all the tables in the namespace.

```
REVOKE READ_IN_NAMESPACE ON NAMESPACE ns1 FROM Kate
```

Here, ns1 is the namespace and Kate is the user.

Note

The label `MODIFY_IN_NAMESPACE` can be used as a helper for granting or revoking permissions to all DDL privileges for a specific namespace to a user or role.

The following example shows:

1. Creation of a namespace and a table.
2. Revocation of the privilege to create any other new tables in the namespace, but allow the table to be dropped.

Example: Namespace Scoped Privileges

```
CREATE NAMESPACE IF NOT EXISTS ns1
```

```
GRANT MODIFY_IN_NAMESPACE ON NAMESPACE ns1 TO usersRole
```

```
CREATE TABLE ns1:t (id INTEGER, name STRING, primary key (id))
```

```
INSERT INTO ns1:t VALUES (1, 'Smith')
```

```
SELECT * FROM ns1:t
```

```
REVOKE CREATE_TABLE_IN_NAMESPACE ON NAMESPACE ns1 FROM usersRole
```

```
DROP NAMESPACE ns1 CASCADE
```

Note

You can save all of the above commands as a **sql** script and execute it in a single command. If you want to execute any of the above commands outside of a SQL prompt, remove the semi colon at the end.

4

Region Management

Oracle NoSQL Database supports Multi-Region Architecture in which you can create tables in multiple KVStores, and still maintain consistent data across these clusters. Each KVStore cluster in a Multi-Region NoSQL Database setup is called a Region. This chapter describes creating and managing regions in Oracle NoSQL Database.

This chapter contains the following topics:

- [CREATE REGION Statement](#)
- [SHOW REGIONS Statement](#)
- [DROP REGION Statement](#)

CREATE REGION Statement

In a Multi-Region Oracle NoSQL Database setup, you must define all the remote regions for each local region. For example, if there are three regions in a Multi-Region setup, you must define the other two regions from each participating region. You use the create region statement to define remote regions in the Multi-Region Oracle NoSQL Database.

Syntax

```
create_region_statement ::= CREATE REGION region_name
```

```
region_name ::= id | DSTRING
```

Semantics

region_name

The name of the region that is different from the local region where the command is executed.

Example 4-1 Create Region

The following create region statement creates a remote region named *my_region1*.

```
CREATE REGION my_region1
```

SET LOCAL REGION Statement

Specifies which region your operations will target in a multi-region environment. By setting a local region, you ensure that your queries and updates are directed to the right regional data store, supporting performance and data locality requirements.

Use the following command to set the local region

```
SET LOCAL REGION my_local_region;
```

Note

You can set the local region only once. After you set the local region, any attempt to change it again results in the following error.

```
sql-> set local region IAD;  
Error handling command set local region IAD:  
Failure 1: 1/SetLocalRegionName IAD failed.: Cannot change local  
region name, multi-region tables are active (25.3.21) on [2026-01-14  
21:03:15.493 UTC]
```

SHOW REGIONS Statement

Syntax

```
show_regions_statement ::= SHOW [AS JSON] REGIONS
```

Semantics

The show regions statement provides the list of regions present in the Multi-Region Oracle NoSQL Database.

AS JSON can be specified if you want the output to be in JSON format.

Example 4-2 Show Regions

The following statement lists all the existing regions.

```
SHOW REGIONS
```

Output:

```
regions  
my_region1 (remote, active)  
my_region2 (remote, active)
```

Example 4-3 Show Regions

The following statement lists all the existing regions in JSON format.

```
SHOW AS JSON REGIONS
```

Output:

```
{"regions" : [  
  {"name" : "my_region1", "type" : "remote", "state" : "active"},  
  {"name" : "my_region2", "type" : "remote", "state" : "active"}  
]}
```

DROP REGION Statement

In a Multi-Region Oracle NoSQL Database environment, the drop region statement removes the specified remote region from the local region.

Syntax

```
drop_region_statement :: DROP REGION region\_name
```

Semantics

region_name

The name of the region that you want to drop. This region must be different from the local region where the command is executed.

Example 4-4 Drop Region

The following drop region statement removes a remote region named *my_region1*.

```
DROP REGION my_region1
```

5

Table Management

In Oracle NoSQL Database, data is stored and organized in tables. This chapter describes tables and creating and managing tables in Oracle NoSQL Database.

A table is an unordered collection of record items, all of which have the same record type. We call this record type the table schema. The table schema is defined by the CREATE TABLE statement. The records of a table are called rows and the record fields are called columns. Therefore, an Oracle NoSQL Database table is a generalization of the (normalized) relational tables found in more traditional RDBMSs.

Although table rows are records, records are not rows. This is because, rows have some additional properties that are not part of the table schema (i.e., they are not stored as top-level columns or nested fields). To extract the values of such properties, the functions listed in the [Functions on Rows](#) section must be used.

This chapter contains the following topics:

- [CREATE TABLE Statement](#)
- [SHOW TABLES Statement](#)
- [DESCRIBE TABLE Statement](#)
- [Table Hierarchies](#)
- [Using the IDENTITY Column](#)
- [Sequence Generator](#)
- [DROP TABLE Statement](#)
- [ALTER TABLE Statement](#)
- [Altering an IDENTITY Column](#)

CREATE TABLE Statement

The table is the basic structure to hold user data. You use the create table statement to create a new table in the Oracle NoSQL Database.

Syntax

```
create_table_statement ::=
    CREATE TABLE [IF NOT EXISTS] table_name [comment]
    "(" table_definition ")" [ttl_definition] [before_image_definition]
    [json_collection_definition]

table_name ::= [namespace\_name ":" ] name_path
name_path ::= field_name ( "." field_name ) *
field_name ::= id | DSTRING

table_definition ::=
    (column_definition | key_definition)
    ( "," (column_definition | key_definition) ) *
```

```

column_definition ::=
    id type\_definition
    [default\_definition | identity\_definition |
uuid\_definition | mr\_counter\_definition]
    [comment]
key_definition ::=
    PRIMARY KEY
    "(" [shard\_key\_definition [","] [id_list_with_size] ")"
    [ttl\_definition]
id_list_with_size ::= id_with_size ("," id_with_size)*
id_with_size ::= id [storage\_size]
storage_size ::= "(" INT_CONSTANT ")"
shard_key_definition ::= SHARD "(" id_list_with_size ")"
ttl_definition ::= USING TTL INT_CONSTANT (HOURS | DAYS)
region_definition ::= IN REGIONS region-name-1,region-name-2 [,...]
before_image_definition ::= ENABLE BEFORE IMAGE [ttl\_definition]
json_collection_definition ::= AS JSON COLLECTION

```

Semantics

table_name

The table name is specified as an optional namespace_name and a local_name. The local name is a name_path because, in the case of child tables, it will consist of a list of dot-separated ids. Child tables are described in the [Table Hierarchies](#) section. A table_name that includes a namespace_name is called a **qualified table name**. When an SQL statement (DDL or DML) references a table by its local name only, the local name is resolved internally to a qualified name with a specific namespace name. See the [Namespace Management](#) chapter.

IF NOT EXISTS

This is an optional clause. If this clause is specified and if a table with the same qualified name exists (or is being created) and if that existing table has the same structure as in the statement, no error is generated. In all other cases and if a table with the same qualified name exists, the create table statement generates an error indicating that the table exists.

ttl_definition

The Time-To-Live (TTL) value is used in computing the expiration time of a row. Expired rows are not included in query results and are eventually removed from the table automatically by Oracle NoSQL Database. If you specify a TTL value while creating the table, it applies as the default TTL for every row inserted into this table. However, you can override the table level TTL by specifying a TTL value via the table insertion API.

The expiration time of a row is computed by adding the TTL value to the current timestamp. To be more specific, for a TTL value of N hours/days, the expiration time is the current time (in UTC) plus N hours/days, rounded up to the next full hour/day. For example, if the current timestamp is 2020-06-23T10:01:36.096 and the TTL is 4 days, the expiration time will be 2020-06-28T00:00:00.000. You can use zero as a special value to indicate that a rows should never expire. If the CREATE TABLE statement has no TTL specification, the default table TTL is zero.

In case of MR Tables with TTL value defined, the rows replicated to other regions carry the expiration time when the row was written. This can be either the default table level TTL value or a row level override that is set by your application. Therefore, this row will expire in all the regions at the same time, irrespective of when they were replicated. However, if a row is updated in one of the regions and it expires in the local region even before it is replicated to one of the remote region(s), then this row will expire as soon as it is replicated and committed in that remote region.

json_collection_definition

The `json_collection_definition` declares the table as a collection of documents. A JSON collection table is a convenient way to store, update, and query your documents. A JSON collection table must include a primary key while creating the table. You can create a JSON collection table with `MR_COUNTERS` if the table is intended to be a multi-region table. For more details on JSON collection table, see [Using JSON Collection Tables](#).

region_definition

This is an optional clause. In case, the table being created is an MR Table, this parameter lists all the regions that the table should span. You must mention at least one remote region in this clause to create the table as an MR Table. For information on MR Tables, see Life Cycle of MR Tables.

Note

Specifying this clause while creating a child table of a MR Table will result in an error.

table_definition

The `table_definition` part of the statement must include at least one field definition, and exactly one primary key definition (Although the syntax allows for multiple `key_definitions`, the query processor enforces the one `key_definition` rule. The syntax is this way to allow for the key definition to appear anywhere among the field definitions).

column_definition

The syntax for a column definition is similar to the `field_definition` grammar rule that defines the fields of a record type. See [Data Type Definitions](#) section. It specifies the name of the column, its data type, whether the column is nullable or not, an optional default value or whether the column is an `IDENTITY` column or not, and an optional comment. As mentioned in [Table Management](#) section, tables are containers of records, and the `table_definitions` acts as an implicit definition of a record type (the table schema), whose fields are defined by the listed `column_definitions`. However, when the `type_definition` grammar rule is used in any DDL statement, the only wildcard type that is allowed is the JSON type. So, for example, it is possible to create a table with a column whose type is JSON, but not a column whose type is ANY.

identity_definition

The `identity_definition` specifies the name of the identity column. There can only be one identity column per table. See [Using the IDENTITY Column](#) section.

uuid_definition

The `uuid_definition` declares the type of a column to be the UUID type. See [Using the UUID data type](#) section.

mr_counter_definition

The `mr_counter_definition` parameter declares the type of a column to be the `MR_COUNTER` datatype. This data type can be used only in a multi-region table. See [Using CRDT datatype](#) in a multi-region table.

key_definition

The syntax for the primary key specification (`key_definition`) specifies the primary key columns of the table as an ordered list of field names. The column names must be among the ones appearing in the `field_definitions`, and their associated type must be one of the following: `INTEGER`, `LONG`, `FLOAT`, `DOUBLE`, `NUMBER`, `STRING`, `ENUM`, `BOOLEAN`, or `TIMESTAMP`. For details, see [Atomic Data Types](#).

The primary key value for each row in a table must be unique.

While Oracle NoSQL Database provides flexibility with respect to data types for the primary key, understanding the tradeoffs associated with different data types is extremely useful when creating your table.

FLOAT and DOUBLE Data type:

FLOAT and DOUBLE data types represent approximate numeric data values. When serialized as a primary key, the DOUBLE data type uses 10 bytes and stores up to 16 decimal places. The FLOAT data type uses 5 bytes and stores up to 8 decimal places. Therefore, DOUBLE data type is more suitable for a primary key as it provides a higher precision and wider range as compared to the FLOAT data type.

ENUM and STRING Data type:

STRING and ENUM represent string data types. The ENUM data type mandates choosing from a list of permitted values. When used as a primary key, the cardinality is small. Also, any addition or deletion to the ENUM values requires schema evolution. The primary keys can only be changed by dropping and recreating the table. Therefore, changing the ENUM values of a primary key is an invasive operation. The STRING data type is better suited for primary keys as its value is a sequence of zero or more Unicode characters.

shard_key_definition

A `key_definition` specifies the table's shard key columns as well, as the first N primary-key columns, where $0 < N \leq M$ and M is the number of primary-key columns. Specification of a shard key is optional. By default, for a root table (a table without a parent) the shard key is the whole primary key. Semantically, the shard key is used to distribute table rows across the multiple servers and processes that comprise an Oracle NoSQL Database store. Briefly, two rows having the same shard key, i.e., the same values on their shardkey columns, will always be located in the same server and managed by the same process. Further details about the distribution of data in Oracle NoSQL Database can be found in the Primary and Shard Key Design section.

storage_size

An additional property of INTEGER-typed primary-key fields is their storage size. This is specified as an integer number between 1 and 5 (the syntax allows any integer, but the query processor enforces the restriction). The storage size specifies the maximum number of bytes that may be used to store in serialized form a value of the associated primary key column. If a value cannot be serialized into the specified number of bytes (or less), an error will be thrown. An internal encoding is used to store INTEGER (and LONG) primary-key values, so that such values are sortable as strings (this is because primary key values are always stored as keys of the "primary" Btree index). The following table shows the range of positive values that can be stored for each byte-size (the ranges are the same for negative values). Users can save storage space by specifying a storage size less than 5, if they know that the key values will be less or equal to the upper bound of the range associated with the chosen storage size.

comment

Comments are included at table-level and they become part of the table's metadata as uninterpreted text. Comments are displayed in the output of the describe statement.

before_image_definition

The `ENABLE BEFORE IMAGE` clause for a table enables the generation and persistent storage of before-images for any subsequent write operations on the table. The before-image is the table row before it gets updated or deleted by a DML operation. You use this clause if your application would like to receive before-images in addition to the after-images in a subscription stream. Streaming before-images enable applications to compute the change or delta made by write operations to tables in a data store. For more details on Oracle NoSQL Database Streams API, see Streams Developer's Guide.

You can include the `ENABLE BEFORE IMAGE` clause with an optional TTL definition to set the TTL value for generated before-images. The before-images TTL determines if the Streams

API can see before-images in the stream. For details, see [Enabling Before-images During Table Creation](#).

Tables used in the examples

This section shows how to create table from the SQL shell for a few sample applications.

Start your KVSTORE or KVLite and open the SQL shell.

```
java -jar lib/kvstore.jar kvlite -secure-config disable
java -jar lib/sql.jar -helper-hosts localhost:5000 -store kvstore
```

See [Running the SQL Shell](#) for the general usage and command options.

Example 5-1 User data application table

The following create table statement defines a users table that holds information about the users:

```
CREATE TABLE users( id INTEGER,
                    firstName STRING,
                    lastName STRING,
                    otherNames ARRAY(RECORD(first STRING, last STRING)),
                    age INTEGER,
                    income INTEGER,
                    address JSON,
                    connections ARRAY(INTEGER),
                    hobbies ARRAY(STRING),
PRIMARY KEY (id) )
```

The rows of the users table defined above represent information about users. For each such user, the otherNames field is an array of records, where each record in the array includes the first and last name fields. The connections field is an array of integers containing the ids of other users that this user is connected with. You can assume that the ids in the array are sorted by some measure of the strength of the connection. The hobbies field is an array of string containing the user's interests in various activities. The address column is a schema-less JSON field.

For example, a typical value for the address field can include the following attributes:

```
{
  "street" : "Pacific Ave",
  "number" : 101,
  "city" : "Santa Cruz",
  "state" : "CA",
  "zip" : 95008,
  "phones" : [
    { "area" : 408, "number" : 4538955, "kind" : "work" },
    { "area" : 831, "number" : 7533341, "kind" : "home" }
  ]
}
```

You can add any number of attributes to the JSON field. The JSON field does not have a fixed data type.

Some addresses may have additional fields, or missing fields, or fields spelled differently. The phones field may not be an array of JSON objects but a single such object. The whole address may be just one string, number, or JNULL.

For more details on the supported data types, see [Oracle NoSQL Database Data Model](#).

Example 5-2 Airline baggage tracking application

The following CREATE TABLE statement defines a BaggageInfo table that holds the checked baggage information of passengers in an airline system:

```
CREATE TABLE BaggageInfo (ticketNo LONG,  
fullName STRING,  
gender STRING,  
contactPhone STRING,  
confNo STRING,  
bagInfo JSON,  
PRIMARY KEY (ticketNo)  
) USING TTL 5 DAYS
```

You can use the baggage tracking schema in an airline's mobile application. The passengers traveling on a flight can track where their checked-in bags or luggage are, along the route to the destination. The mobile application uses the Oracle NoSQL Database to store all the data related to the baggage. In the backend, the mobile application logic uses SQL queries to retrieve the required data.

This baggage tracking schema creates a BaggageInfo table with columns that include atomic data types and JSON data type to hold passenger information. The BaggageInfo table has the passenger's ticket number as the primary key column. The fullName, gender, contactPhone, and confNo (reservation number) fields store the passenger's information, which is part of a fixed schema. This table also includes a bagInfo column as a schema-less JSON array, which represents the tracking information of a passenger's checked-in baggage. In contrast to the fixed schema fields, you can add any number of attributes to this schema-less JSON field.

See [Example 5-5](#) for a sample row from the BaggageInfo table.

You can create the table by specifying a TTL value, after which the rows expire automatically and are not available anymore. The TTL value is followed by either HOURS or DAYS. In this schema, the rows of the table expire after 5 days. You can check the hours remaining until a row expires using the remaining_days [Built-in Functions](#).

Specifying the TTL value while creating a table is optional. If you don't provide a TTL value, the rows of a table will have an infinite expiration time.

Example 5-3 Streaming media service

The following CREATE TABLE statement defines a stream_acct table that creates a TV streaming application:

```
CREATE TABLE stream_acct(  
acct_id INTEGER,  
profile_name STRING,  
acct_data JSON,  
PRIMARY KEY(acct_id)  
)
```

The Streaming Media Service application streams various shows across the globe. Every show has several seasons, and every season has multiple episodes. You need a persistent metadata store that keeps track of the current activity of the subscriber using the TV streaming application.

You can use the TV streaming application to provide useful information to the subscriber such as episodes they watched, watch time per episode, total number of seasons of the show they watched, and so forth. The mobile application uses the Oracle NoSQL Database to store the data and performs SQL queries to retrieve the required data and offer it to the user.

The Streaming Media Service schema creates a `stream_acct` table with columns that include atomic data types and JSON data type to hold subscriber details.

The `stream_acct` table has the subscriber's account ID as the primary key column. The other fixed schema fields `profile_name` and `account_expiry` contain the subscriber's details. The `stream_acct` table includes an `acct_data` column as a schema-less JSON field to keep track of the subscriber's current activity. In contrast to the fixed schema fields, you can add any number of attributes to this schema-less JSON field.

See [Example 5-6](#) for a sample row from the `stream_acct` table.

Load data to the tables

You can load data into the tables using [INSERT](#) statement from the SQL shell directly by supplying the statement in the SQL prompt or by using scripts.

Example 5-4 Insert the data into the User data application (users table)

The following example shows how to insert the data into the User data - users table using the INSERT statement directly in the SQL prompt:

```
INSERT INTO users VALUES (
    10,
    "John",
    "Smith",
    [ {"first" : "Johnny", "last" : "Good"}, {"first" : "Johnny2", "last" :
"Brave"}, {"first" : "Johnny3", "last" : "Kind"}, {"first" : "Johnny4", "last" :
"Humble"} ],
    22,
    45000,
    {
        "street" : "Pacific Ave",
        "number" : 101,
        "city" : "Santa Cruz",
        "state" : "CA",
        "zip" : 95008,
        "phones" : [
            { "area" : 408, "number" : 4538955, "kind" : "work" },
            { "area" : 831, "number" : 7533341, "kind" : "home" },
            { "area" : 831, "number" : 7533382, "kind" : "mobile" }
        ]
    },
    [ 30, 55, 43 ],
    [ "Reading", "Soccer", "Hiking", "Gardening" ]
)
```

```
INSERT INTO users VALUES (
```

```

20,
"Jane",
"Smith",
[ {"first" : "Jane", "last" : "Charming"} ],
22,
55000,
{
  "street" : "Atlantic Ave",
  "number" : 201,
  "city" : "San Jose",
  "state" : "CA",
  "zip" : 95005,
  "phones" : [
    { "area" : 608, "number" : 6538955, "kind" : "work" },
    { "area" : 931, "number" : 9533341, "kind" : "home" },
    { "area" : 931, "number" : 9533382, "kind" : "mobile" }
  ]
},
[ 40, 75, 63 ],
[ "Knitting", "Hiking", "Baking", "BingeWatching" ]
)

INSERT INTO users VALUES (

30,
"Adam",
"Smith",
[ {"first" : "Adam", "last" : "BeGood"} ],
45,
75000,
{
  "street" : "Indian Ave",
  "number" : 301,
  "city" : "Houston",
  "state" : "TX",
  "zip" : 95075,
  "phones" : [
    { "area" : 618, "number" : 6618955, "kind" : "work" },
    { "area" : 951, "number" : 9613341, "kind" : "home" },
    { "area" : 981, "number" : 9613382, "kind" : "mobile" }
  ]
},
[ 60, 45, 73 ],
[ "Soccer", "Riding", "PianoForte", "RockClimbing", "Sketching" ]
)

INSERT INTO users VALUES (

40,
"Joanna",
"Smith",
[ {"first" : "Joanna", "last" : "Smart"} ],
NULL,
75000,
{

```

```

        "street" : "Tex Ave",
        "number" : 401,
        "city" : "Houston",
        "state" : "TX",
        "zip" : 95085,
        "phones" : [
            { "area" : NULL, "number" : 1618955, "kind" : "work" },
            { "area" : 451, "number" : 4613341, "kind" : "home" },
            { "area" : 481, "number" : 4613382, "kind" : "mobile" }
        ]
    },
    [ 70, 30, 40 ],
    ["Soccer", "Riding", "PianoForte", "RockClimbing", "Sketching"]
)

```

Example 5-5 Import data into airline baggage tracking application (BaggageInfo) table

The following example shows how to create sample table for the airline baggage tracking application and import data using script from the SQL prompt:

Download the script `baggageschema_loaddata.sql`.

Using the `load` command, run the required script.

```
load -file baggageschema_loaddata.sql
```

This creates the `BaggageInfo` table used in the examples and loads data into the table.

One sample row from the airline baggage tracking application - `BaggageInfo` table is shown below.

The passenger's ticket number, `ticketNo` is the primary key of the table. The `fullName`, `gender`, `contactPhone`, and `confNo` (reservation number) fields store the passenger's information, which is part of a fixed schema. The `bagInfo` column is a schema-less JSON array, which represents the tracking information of a passenger's checked-in baggage.

Each element of the `bagInfo` array corresponds to a single checked-in bag. The size of the `bagInfo` array gives the total bags checked-in by a passenger. Each bag has an `id` and a `tagnum` field. The `routing` field includes the routing information from the passenger's travel itinerary. The `lastActionCode` and `lastActionDesc` fields hold the latest action taken on the bag and its action code at the current destination. The `lastSeenStation` field includes the airport code of the bag's current destination. The `lastSeenTimeGmt` field includes the latest action time. The `bagArrivalDate` field holds the expected arrival date at the destination airport. The `bagInfo` array further includes a nested `flightLegs` array with fields to track the source and transit details.

Each element of the `flightLegs` array corresponds to a travel leg. The fields `flightNo` holds the flight number, `flightDate` holds the departure date, `fltRouteSrc` holds the originating airport code, and `fltRouteDest` field hold the destination airport code for each travel leg. The `flightLegs` array further includes a nested `actions` array with fields to track the activities performed on the checked bag at each travel leg.

Each element of the `actions` array includes the fields `actionAt`, `actionCode`, and `actionTime` to track the tasks at source and destination airports in each travel leg.

```

"ticketNo" : 1762344493810,
"fullName" : "Adam Phillips",

```

```

"gender" : "M",
"contactPhone" : "893-324-1064",
"confNo" : "LE6J4Z",
[ {
  "id" : "79039899165297",
  "tagNum" : "17657806255240",
  "routing" : "MIA/LAX/MEL",
  "lastActionCode" : "OFFLOAD",
  "lastActionDesc" : "OFFLOAD",
  "lastSeenStation" : "MEL",
  "flightLegs" : [ {
    "flightNo" : "BM604",
    "flightDate" : "2019-02-01T01:00:00",
    "fltRouteSrc" : "MIA",
    "fltRouteDest" : "LAX",
    "estimatedArrival" : "2019-02-01T03:00:00",
    "actions" : [ {
      "actionAt" : "MIA",
      "actionCode" : "ONLOAD to LAX",
      "actionTime" : "2019-02-01T01:13:00"
    }, {
      "actionAt" : "MIA",
      "actionCode" : "BagTag Scan at MIA",
      "actionTime" : "2019-02-01T00:47:00"
    }, {
      "actionAt" : "MIA",
      "actionCode" : "Checkin at MIA",
      "actionTime" : "2019-02-01T23:38:00"
    } ]
  } ],
}, {
  "flightNo" : "BM667",
  "flightDate" : "2019-01-31T22:13:00",
  "fltRouteSrc" : "LAX",
  "fltRouteDest" : "MEL",
  "estimatedArrival" : "2019-02-02T03:15:00",
  "actions" : [ {
    "actionAt" : "MEL",
    "actionCode" : "Offload to Carousel at MEL",
    "actionTime" : "2019-02-02T03:15:00"
  }, {
    "actionAt" : "LAX",
    "actionCode" : "ONLOAD to MEL",
    "actionTime" : "2019-02-01T07:35:00"
  }, {
    "actionAt" : "LAX",
    "actionCode" : "OFFLOAD from LAX",
    "actionTime" : "2019-02-01T07:18:00"
  } ]
} ],
"lastSeenTimeGmt" : "2019-02-02T03:13:00",
"bagArrivalDate" : "2019.02.02T03:13:00"
} ]

```


Example 5-6 Import data into TV streaming application (stream_acct) table

The following example shows how to create sample table for TV streaming application, and import data using script from the SQL prompt:

Download the script acctstream_loaddata.sql.

Using the load command, run the required script.

```
load -file acctstream_loaddata.sql
```

This creates the stream_acct table used in the examples and loads data into the table.

One sample row from the Streaming Media Service - stream_acct table is shown below.

The subscriber's account ID, acct_id is the primary key of the table. The fields profile_name and account_expiry contain the subscriber's details. The acct_data column is a schema-less JSON field, which keeps track of the subscriber's current activity.

Each element of the acct_data JSON represents a user with the given subscriber's profile name. User data contains the fields firstName, lastName, and country to hold user information. The acct_data JSON field further includes a nested contentStreamed JSON array to track the shows watched by the user.

Each element of the contentStreamed array contains the showName field to store the name of the show. The showId field includes the identifier of the show. The showtype field indicates the type such as tvseries, sitcom, and so forth. The genres array lists the show's categorization. The numSeasons field contains the total number of seasons streamed for the show. The contentStreamed JSON array also includes a nested seriesInfo JSON array to track the watched episodes.

Each element of the seriesInfo array contains a seasonNum field to identify the season. The numEpisodes field indicates the total number of episodes streamed in the given season. The seriesInfo array further includes an episodes array to track the details of each watched episode.

Each element of the episodes array contains the episodeID field to identify the episode. The episodeName field includes the episode's name. The lengthMin field includes the show's telecast duration in minutes. The minWatched field includes the duration for which a user has watched the episode. The date field includes the date on which the user watched the given episode.

```
1,
123456789,
"AP",
"2023-10-18",
{
  "firstName": "Adam",
  "lastName": "Phillips",
  "country": "Germany",
  "contentStreamed": [
    {
      "showName": "At the Ranch",
      "showId": 26,
      "showtype": "tvseries",
      "genres": ["action", "crime", "spanish"],
      "numSeasons": 4,
      "seriesInfo": [
```

```

{
  "seasonNum" : 1,
  "numEpisodes" : 2,
  "episodes": [
    {
      "episodeID": 20,
      "episodeName" : "Season 1 episode 1",
      "lengthMin": 85,
      "minWatched": 85,
      "date" : "2022-04-18"
    },
    {
      "episodeID": 30,
      "lengthMin": 60,
      "episodeName" : "Season 1 episode 2",
      "minWatched": 60,
      "date" : "2022-04-18"
    }
  ]
},
{
  "seasonNum": 2,
  "numEpisodes" : 2,
  "episodes": [
    {
      "episodeID": 40,
      "episodeName" : "Season 2 episode 1",
      "lengthMin": 50,
      "minWatched": 50,
      "date" : "2022-04-25"
    },
    {
      "episodeID": 50,
      "episodeName" : "Season 2 episode 2",
      "lengthMin": 45,
      "minWatched": 30,
      "date" : "2022-04-27"
    }
  ]
}
],
{
  "seasonNum": 3,
  "numEpisodes" : 2,
  "episodes": [
    {
      "episodeID": 60,
      "episodeName" : "Season 3 episode 1",
      "lengthMin": 50,
      "minWatched": 50,
      "date" : "2022-04-25"
    },
    {
      "episodeID": 70,
      "episodeName" : "Season 3 episode 2",

```

```

        "lengthMin": 45,
        "minWatched": 30,
        "date" : "2022-04-27"
    }
  ]
}
},
{
  "showName": "Bienvenu",
  "showId": 15,
  "showtype": "tvseries",
  "genres" : ["comedy", "french"],
  "numSeasons" : 2,
  "seriesInfo": [
    {
      "seasonNum" : 1,
      "numEpisodes" : 2,
      "episodes": [
        {
          "episodeID": 20,
          "episodeName" : "Bonjour",
          "lengthMin": 45,
          "minWatched": 45,
          "date" : "2022-03-07"
        },
        {
          "episodeID": 30,
          "episodeName" : "Merci",
          "lengthMin": 42,
          "minWatched": 42,
          "date" : "2022-03-08"
        }
      ]
    }
  ]
}
]
}
}

```

SHOW TABLES Statement

Syntax

```

show_tables_statement ::=
    SHOW [AS JSON] (TABLES | TABLE table_name)

```

Semantics

The show tables statement provides the list of tables present in the system. If you want to know the details of a specific table, then you can use show table statement. If the named table does not exist then this statement fails.

Example 5-7 Show Tables

The following statement lists all the tables in the system.

```
SHOW TABLES
```

Output:

```
tables
  SYS$IndexStatsLease
  SYS$PartitionStatsLease
  SYS$SGAttributesTable
  SYS$StreamRequest
  SYS$StreamResponse
  SYS$TableStatsIndex
  SYS$TableStatsPartition
  Users2
  users
```

Example 5-8 Show Tables

The following statement lists all the tables in the system in JSON format.

```
SHOW AS JSON TABLES
```

Output:

```
{"tables" : [
  "SYS$IndexStatsLease",
  "SYS$PartitionStatsLease",
  "SYS$SGAttributesTable",
  "SYS$StreamRequest",
  "SYS$StreamResponse",
  "SYS$TableStatsIndex",
  "SYS$TableStatsPartition",
  "Users2",
  "users"
]}
```

Example 5-9 Show Tables

The following statement lists a specific table in the system.

```
SHOW TABLE users
```

Output:

```
tableHierarchy
  users
```

DESCRIBE TABLE Statement

Syntax

```
describe_table_statement ::=
    (DESCRIBE | DESC) [AS JSON] TABLE table_name
    [ "(" field_name ["," field_name] ")"]
```

Semantics

The description for tables contains the following information:

- Name of the table.
- Time-To-Live value of the table.
- Owner of the table.
- Whether the table is a system table.
- Name of parent tables.
- Name of children tables.
- List of indexes present on the table.
- Description of the table.

The description for fields contains the following information:

- Id of the field.
- Name of the field.
- Datatype of fields, for example, INTEGER, STRING, Map(INTEGER), etc.
- Whether the field is nullable. If the field is nullable then 'Y' is displayed, otherwise 'N' is displayed.
- Default value of the field.
- Whether the field is a shard key. If the field is a shard key then 'Y' is displayed, otherwise 'N' is displayed.
- Whether the field is a primary key. If the field is a primary key then 'Y' is displayed, otherwise 'N' is displayed.
- Whether the field is an identity field. If the field is an identity field then 'Y' is displayed, otherwise 'N' is displayed.

AS JSON can be specified if you want the output to be in JSON format.

Example 5-10 Describe Table

AS JSON can be specified if you want the output to be in JSON format.

```
DESCRIBE TABLE users
```

Output:

```

=== Information ===
+-----+-----+-----+-----+-----+-----+-----+-----+

```

```

+-----+
| name | ttl | owner | sysTable | r2compat | parent | children | indexes |
description |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| users |      |      | N      | N      |      |      |      |
|      |      |      |      |      |      |      |      |
+-----+-----+-----+-----+-----+-----+-----+
+-----+

=== Fields ===
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| id | name | type | nullable | default | shardKey |
primaryKey | identity |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 1 | id | Integer | N | NullValue | Y |
Y |      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 2 | firstName | String | Y | NullValue |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 3 | lastName | String | Y | NullValue |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 4 | otherNames | Array( | Y | NullValue |
|      |      |      |      |      |
|      |      | RECORD( |      |      |
|      |      | first : String, |      |      |
|      |      | last : String |      |      |
|      |      | )) |      |      |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 5 | age | Integer | Y | NullValue |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 6 | income | Integer | Y | NullValue |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 7 | address | Json | Y | NullValue |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 8 | connections | Array(Integer) | Y | NullValue |
|      |      |      |      |      |
+---+-----+-----+-----+-----+-----+-----+

```

```

+-----+-----+
| 9 | expenses | Map(Integer) | Y | NullValue |
|   |           |              |   |           |
+---+-----+-----+-----+-----+-----+-----+
+-----+-----+

```

Example 5-11 Describe Table

The following statement provides information about the users table and its fields in JSON format.

```
DESC AS JSON TABLE users
```

Output:

```

{
  "json_version" : 2,
  "type" : "table",
  "name" : "users",
  "id" : "57",
  "fields" : [{
    "name" : "id",
    "type" : "INTEGER",
    "nullable" : false
  }, {
    "name" : "firstName",
    "type" : "STRING",
    "nullable" : true
  }, {
    "name" : "lastName",
    "type" : "STRING",
    "nullable" : true
  }, {
    "name" : "otherNames",
    "type" : "ARRAY",
    "nullable" : true,
    "collection" : {
      "type" : "RECORD",
      "fields" : [{
        "name" : "first",
        "type" : "STRING",
        "nullable" : true
      }, {
        "name" : "last",
        "type" : "STRING",
        "nullable" : true
      }]
    }
  }, {
    "name" : "age",
    "type" : "INTEGER",
    "nullable" : true
  }, {
    "name" : "income",
    "type" : "INTEGER",

```

```

        "nullable" : true
    }, {
        "name" : "address",
        "type" : "JSON",
        "nullable" : true
    }, {
        "name" : "connections",
        "type" : "ARRAY",
        "nullable" : true,
        "collection" : {
            "type" : "INTEGER"
        }
    }, {
        "name" : "hobbies",
        "type" : "ARRAY",
        "nullable" : true,
        "collection" : {
            "type" : "STRING"
        }
    }
}],
"primaryKey" : ["id"],
"shardKey" : ["id"],
"indexes" : [{
    "name" : "idx_modtime",
    "type" : "secondary",
    "fields" : [{
        "function" : "modification_time",
        "type" : "Timestamp(3)"
    }
}],
"withNoNulls" : false,
"withUniqueKeysPerRow" : false
}]
}

```

Example 5-12 Describe Table

The following statement provides information about a specific field in the users table.

```
DESCRIBE TABLE users (income)
```

Output:


```

+---+-----+-----+-----+-----+-----+-----+
+-----+
| id | name | type  | nullable | default | shardKey | primaryKey |
identity |
+---+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | income | Integer | Y        | NullValue |          |            |
|          |
+---+-----+-----+-----+-----+-----+-----+
+-----+

```


Table Hierarchies

The Oracle NoSQL Database enables tables to exist in a parent-child relationship. This is known as table hierarchies.



Terminologies used in a table hierarchy:

- **Root Node:** The topmost node of a tree or the node that does not have any parent node is called the root node. In the above diagram, A is the root node.
- **Parent Node:** The immediate predecessors of a node are called its parent nodes. For example, in the above diagram, A is the parent of B and G.
- **Child Node:** The immediate successors of a node are called its child nodes. For example, in the above diagram, B and G are children of node A.
- **Leaf Node:** The node that does not have any child node is called a leaf node. In the above diagram, C, D, H and I are leaf nodes.
- **Ancestor:** All the predecessor nodes on the paths from the root to that node are called its ancestor nodes. For example, in the above diagram, A and G are ancestors of H and I. A is the ancestor of G, H and I.
- **Descendant:** All the successor nodes on the paths from that node to the leaf nodes are called its descendant nodes. For example, in the above diagram, H and I are descendants of A and G. G, H and I are descendants of A.
- **Sibling:** Children of the same parent node are called siblings. For example, in the above diagram, B and G are siblings, as they are both children of node A. Similarly H and I are also siblings as they are children of node G.

Why do you need hierarchical tables in Oracle NoSQL Database?

Given that the Oracle NoSQL Database model includes arrays and maps, one may wonder why child tables are needed? Creating child tables is a way to achieve hierarchical data modeling. It optimizes scalability and performance as well as provides a way to easily perform scalable ACID transactions on collections of records.

1. For each parent row, it is possible to store matching child records inside an array or map within the parent row. However, this can lead to very large parent rows, resulting in degraded performance, especially given the append-only architecture of Oracle NoSQL Database, where each update creates a new version of the entire row. Child tables are recommended when a parent contains many/large child records, as they provide a scalable way to represent one-to-many or hierarchical relationships.

2. In Oracle NoSQL Database, ACID transactions (guaranteeing Atomicity, Consistency, Isolation, and Durability) are supported for records that are co-located on the same shard (that is, shard key values are the same). Hierarchical tables inherently achieve this co-location property as the child records inherit the parent table's primary key, ensuring all records in the hierarchy are physically co-located on the same shard by design.

Note

The shard key is always a subset of the primary key.

3. By structuring your data as hierarchical tables, you can perform atomic transactions on the parent and any/all of its children together, maintaining strong data integrity for the entire hierarchy within Oracle NoSQL Database, while not sacrificing the scalability.
4. Hierarchical modeling ensures related data is managed together. It allows efficient retrieval of all child records for a parent using a single query, eliminating the need for multiple separate queries. For further details on the supported queries for hierarchical tables, see [Joins](#).

Hierarchical tables in an Oracle NoSQL Database are :


- Very efficient for write-heavy workloads.
- More flexible for fine-grained authorization. Authorization is permission given to a user to access a resource. In a fine-grained authorization, the access rights given to a user for a resource may vary by conditions at run-time. In a hierarchical setup, access rights given to the parent table might be different from the access rights given to the child table, and so it is more flexible.

How do hierarchical tables work in an Oracle NoSQL Database?

The data store's replication nodes are organized into shards. A single shard contains multiple replication nodes and a master node. A shard key is created to distribute data across the Oracle NoSQL Database cluster for scalability. Records with the same shard key are co-located for easy reference and access. In a hierarchical table, the child table inherits the primary key columns, which contains the shard key columns of its parent table. This is done implicitly, without including the parent columns in the CREATE TABLE statement of the child. All tables in the hierarchy have the same shard key columns.

Creating hierarchical tables

The create table statement allows for a table to be created as a child of another table, which then becomes the parent of the new table. This is done by using a composite name (a name_path) for the child table. A composite name consists of a number N ($N > 1$) of identifiers separated by dots. The last identifier is the *local* name of the child table and the first N-1 identifiers are the name of the parent.



Semantics

The semantic implications of a parent-child relationship are the following:

- A child table inherits the primary key columns of its parent table. This is done implicitly, without including the parent columns in the create table statement of the child. For example, in the following [Example 5-13](#) example, table A.B has an extra column, called ida, and its primary key columns are ida and idb. Similarly, table A.B.C has 2 extra columns, ida and idb, and its primary key columns are ida, idb, and idc. The inherited columns are placed first in the schema of a child table.
- All tables in the hierarchy have the same shard key columns, which are specified in the create table statement of the root table. So, in our example, the common shard key is column ida. Trying to include a shard key clause in the create table statement of a non-root table will raise an error.
- A parent table cannot be dropped before its children are dropped.
- When two rows RC and RP from a child table C and its parent table P, respectively, have the same values on their common primary key columns, we say that RP and RC match, or that RP contains RC. In this case, RP and RC will also be co-located physically, because they have the same shard key. Given that a child table always has more primary key columns than its parent, a parent row may contain multiple child rows, but a child row will match with at most one parent row.

Note

Oracle NoSQL Database does not require that all the rows in a child table have a matching row in the parent table. In other words, a referential integrity constraint is not enforced.

Example 5-13 Table Hierarchy

The following statements create a table hierarchy, that is a tree of tables connected by parent-child relationships. A is the root table, A.B and A.G are children of A, and A.B.C is a child of A.B (and a grandchild of A).

```
CREATE TABLE A (ida INTEGER, a1 STRING, a2 INTEGER, PRIMARY KEY(ida))
```

```
CREATE TABLE A.B (idb INTEGER, b1 STRING, a2 STRING, PRIMARY KEY(idb))
```

```
CREATE TABLE A.B.C (idc INTEGER, b1 STRING, c2 STRING, PRIMARY KEY(idc))
```

```
CREATE TABLE A.G (idg INTEGER, g1 STRING, g2 DOUBLE, PRIMARY KEY(idg))
```

Table Hierarchy in a Multi-Region table:

You can create child tables in an existing Multi-Region architecture.

Example 5-14 Table Hierarchy in a Multi-Region table - create table in two regions

Create the table `users` in two regions, `FRA` and `LON`.

```
CREATE TABLE users (  
    id INTEGER,  
    name STRING,  
    team STRING,  
    PRIMARY KEY (id))  
IN REGIONS FRA,LON
```

Under the `users` table, you can create a child table using this statement.

```
CREATE TABLE users.userdet (  
    pan INTEGER,  
    address STRING,  
    email STRING,  
    PRIMARY KEY(pan))
```

Specifying the `REGIONS` clause while creating a Multi-Region child table will result in an error as illustrated below.

```
CREATE TABLE users.userinfo (pan INTEGER, address STRING, email STRING,  
PRIMARY KEY(pan) IN REGIONS FRA,LON)
```

Output:

```
Error handling command CREATE TABLE users.userinfo (  
    pan INTEGER,  
    address STRING,  
    email STRING,  
    PRIMARY KEY(pan) IN REGIONS FRA,LON): Error: at (5, 24) missing ')' at  
'IN', at line 5:24  
rule stack: [parse, statement, create_table_statement]
```

You can view the description of the Multi-Region child table as shown below. Note that the child table automatically inherits the primary key columns of its parent table.

```
desc as json table users.userdet
```

```
{  
  "json_version" : 2,  
  "type" : "table",  
  "name" : "userdet",  
  "id" : "63",  
  "parent" : "users",  
  "regions" : {  
    "2" : "FRA",  
    "1" : "LON"  
  },  
  "fields" : [{  
    "name" : "id",  
    "type" : "INTEGER",
```

```

        "nullable" : false
    }, {
        "name" : "pan",
        "type" : "INTEGER",
        "nullable" : false
    }, {
        "name" : "address",
        "type" : "STRING",
        "nullable" : true
    }, {
        "name" : "email",
        "type" : "STRING",
        "nullable" : true
    }
    ],
    "primaryKey" : ["id", "pan"],
    "shardKey" : ["id"]
}

```

Example: Hierarchical Table

This section creates the hierarchical tables and loads the data for a sample baggage tracking application.

Airline Baggage Tracking Application Using Hierarchical Tables

Start your KVSTORE or KVLite and open the SQL shell.

```
java -jar lib/kvstore.jar kvlite -secure-config disable
```

```
java -jar lib/sql.jar -helper-hosts localhost:5000 -store kvstore
```

See [Running the SQL Shell](#) for the general usage and command options.

The baggage tracking application that we saw in [Example 5-2](#) can also be realized using a parent child schema. For every flight ticket number, there is a passenger and their baggage associated with it. The root table is `ticket`, and it has 2 child tables `passengerInfo` and `baggageInfo`. The `passengerInfo` table contains the details of the passenger and the `baggageInfo` contains details of the bags checked in by the passenger. These bags are tracked through their transit through multiple intermediary stations. This tracking information is captured in a table called `flight legs` which is the child of the `baggageInfo` table. Following are the tables created:

```

### CREATE table ticket if not present ###
CREATE TABLE IF NOT EXISTS ticket(
    ticketNo LONG,
    confNo STRING,
    PRIMARY KEY(ticketNo))

### CREATE table ticket.baginfo if not present ###
CREATE TABLE IF NOT EXISTS ticket.bagInfo(
    id LONG,
    tagNum LONG,
    routing STRING,
    lastActionCode STRING,

```

```
lastActionDesc STRING,  
lastSeenStation STRING,  
lastSeenTimeGmt TIMESTAMP(4),  
bagArrivalDate TIMESTAMP(4),  
PRIMARY KEY(id))  
  
### CREATE table ticket.bagInfo.flightLegs if not present ###  
CREATE TABLE IF NOT EXISTS ticket.bagInfo.flightLegs(  
flightNo STRING,  
flightDate TIMESTAMP(4),  
fltRouteSrc STRING,  
fltRouteDest STRING,  
estimatedArrival TIMESTAMP(4),  
actions JSON, PRIMARY KEY(flightNo))  
  
### CREATE table ticket.passengerInfo if not present ###  
CREATE TABLE IF NOT EXISTS ticket.passengerInfo(  
contactPhone STRING,  
fullName STRING,  
gender STRING,  
PRIMARY KEY(contactPhone))
```

Let us see a sample row from each of the tables:

Ticket:

```
{  
  "ticketNo": 1762344493810,  
  "id": 79039899165297,  
  "tagNum": 17657806255240,  
  "routing": "MIA/LAX/MEL",  
  "lastActionCode": "OFFLOAD",  
  "lastActionDesc": "OFFLOAD",  
  "lastSeenStation": "MEL",  
  "lastSeenTimeGmt": "2019-02-01T16:13:00Z",  
  "bagArrivalDate": "2019-02-01T16:13:00Z",  
  "actions": null  
}
```

Passenger Info:

```
{  
  "ticketNo": 1762344493810,  
  "contactPhone": "893-324-1064",  
  "fullName": "Adam Phillips",  
  "gender": "M"  
},
```

Baggage Info:

```
{  
  "ticketNo": 1762344493810,  
  "id": 79039899165297,  
  "tagNum": 17657806255240,  
  "routing": "MIA/LAX/MEL",  
  "lastActionCode": "OFFLOAD",  
  "lastActionDesc": "OFFLOAD",  
  "lastSeenStation": "MEL",  
  "lastSeenTimeGmt": "2019-02-01T16:13:00Z",
```

```

        "bagArrivalDate": "2019-02-01T16:13:00Z",
        "actions": null
    },
    Flight Legs:
    {
        "ticketNo": 1762344493810,
        "id": 79039899165297,
        "flightNo": "BM604",
        "flightDate": "2019-02-01T06:00:00Z",
        "fltRouteSrc": "MIA",
        "fltRouteDest": "LAX",
        "estimatedArrival": "2019-02-01T11:00:00Z",
        "actions": [
            {
                "actionAt": "MIA",
                "actionCode": "ONLOAD to LAX",
                "actionTime": "2019-02-01T06:13:00Z"
            },
            {
                "actionAt": "MIA",
                "actionCode": "BagTag Scan at MIA",
                "actionTime": "2019-02-01T05:47:00Z"
            },
            {
                "actionAt": "MIA",
                "actionCode": "Checkin at MIA",
                "actionTime": "2019-02-01T04:38:00Z"
            }
        ]
    }
}

```

Download the script `parentchildtbls_loaddata.sql`

Use the load command to run the script. The tables are created and the data from the JSON files is loaded into the tables.

```
load -file parentchildtbls_loaddata.sql
```

The `parentchildtbls_loaddata.sql` contains the following:

```

### Begin Script ###
load -file parentchild.ddl
import -table ticket -file ticket.json
import -table ticket.bagInfo -file bagInfo.json
import -table ticket.passengerInfo -file passengerInfo.json
import -table ticket.bagInfo.flightLegs -file flightLegs.json
### End Script ###

```

Using JSON Collection Tables

A JSON collection table facilitates the creation of a schema-less table. You can define the [primary key fields](#) using several NoSQL data types (See [Creating JSON collection table](#)). You can add data into the table using an [INSERT](#) or [UPSERT](#) statement. Each row is a single document, which can contain any number of JSON fields.

You can create the JSON collection tables with optional TTL values or add the TTL values later by modifying the table. No other alteration is allowed on JSON collection tables.

You can also create rows in JSON collection table with user-supplied metadata. Using the respective APIs, you can insert, retrieve, update, and delete the metadata. For more information on how to manage user-supplied metadata, see *Using User-Supplied Metadata in Developers Guide*.

Significance of JSON Collection Tables

JSON collection tables are useful for applications that store and retrieve data purely as documents. It is a convenient development paradigm for those applications where using a fixed schema is not optimal. Tables with a fixed schema require the specification of column names during creation and DML operations. The addition or removal of fields from fixed schema tables requires schema evolution, which is a costly operation.

As the JSON collection tables are schema-less, you do not have to modify the schema to add or remove documents from the table. The rows in the JSON collection tables can differ from one another in terms of the number of JSON fields that can be stored without predeclaring the JSON column names.

Creating JSON collection table

You can use the following syntax to create the JSON collection table.

Syntax

```
json_collection_definition ::= AS JSON COLLECTION
```

Semantics

You create a JSON collection table with one or more [primary key fields](#). When you insert data into the table, a single document is created without an explicit declaration of a column of type JSON. The document can contain any number of JSON fields. The JSON fields must be valid [JSON data types](#). Using non-JSON data types while inserting data will cause an exception. The primary key fields can include [IDENTITY columns](#) and [UUIDs](#). The JSON collection tables can't be changed into fixed-schema tables.

You can create a JSON collection table with [MR_COUNTERS](#) if the table is intended to be a multi-region table. The MR_COUNTERS must be declared at the time of table creation and must be only top-level fields in the document. You can't include MR_COUNTERS elsewhere in the document.

You can supply an optional TTL value during the table creation. You can use the [ALTER TABLE statement](#) to only modify the TTL values.

The JSON collection tables support the following data types for the primary key fields:

Table 5-1 Data types for primary key fields in JSON collection tables

Field type	Supported data type
Primary key fields	- integer - long - double - float - number - string

The following examples demonstrate the creation of JSON collection tables. For inserting data into the tables, see [Inserting Rows into JSON Collection Tables](#).

JSON collection table for a sample application

Example 5-15 Create a JSON collection table for a shopping application

```
CREATE TABLE storeAcct(contactPhone string, primary key(contactPhone)) AS
JSON COLLECTION
```

Explanation: In the above DDL statement, you create a JSON collection table for a shopping application. This table includes the contactPhone as the primary key field. There is no need to supply any other field except the primary key field in the DDL command.

When you insert data into the table, the JSON collection table automatically considers the inserted fields other than the primary key fields to be JSON fields.

JSON collection table with more than one primary key fields

Example 5-16 Create a JSON collection table with an IDENTITY column as one of the primary key fields and a TTL value of 90 days

```
CREATE TABLE storeAcctComposite(contactPhone string, id integer generated by
default as
identity, primary key(contactPhone, id)) AS JSON COLLECTION USING TTL
90 DAYS
```

Explanation: The table includes two primary key fields, one of them is a contactPhone field and the other is id, which is declared as an IDENTITY column to autogenerate the order IDs. In this example, you create the table with a [TTL value](#) of 90 days. In the JSON collection DDL, you supply the USING TTL clause followed by the number of hours/days to create with TTL.

JSON collection table with MR_COUNTERS

Example 5-17 Create a JSON collection table for a shopping application with MR_COUNTERS in two regions

```
CREATE TABLE storeAcctMR(contactPhone string, primary key(contactPhone),
mycounter as integer mr_counter) in regions FRA, LON AS JSON COLLECTION
```

Explanation: In the above DDL statement, you create a JSON collection table with an [MR_COUNTER](#) data type in two regions FRA and LON. To create regions, see [CREATE REGION Statement](#).

You must define a table column as an MR_COUNTER along with its subtype during the table creation. When you insert data into this table, excluding primary key fields and MR_COUNTERS, all other supplied fields are implicitly added as JSON fields.

Note

- Declaring MR_COUNTERS in nested JSON fields in the document is not supported and will return an error if tried. The MR_COUNTERS must be top-level fields in the document.

Example 1: MR_COUNTER as a top-level field in the document after inserting data into the storeAcctMR table above.

```
{
  "contactPhone" : "1817113382",
  "address" : {
    "city" : "Houston",
    "number" : 401,
    "state" : "TX",
    "street" : "Tex Ave",
    "zip" : 95085
  },
  "firstName" : "Adam",
  "lastName" : "Smith",
  "mycounter" : 0
}
```

Example 2: In this example, you are trying to declare an MR_COUNTER (counter) in the address field of the document, which is not supported and returns the following error:

```
create table storeAcctMR(contactPhone string, primary
key(contactPhone), mycounter as integer mr_counter, address.counter
as integer mr_counter) in regions FRA, LON as json collection
```

```
Error handling command create table storeAcctMR(contactPhone
string, primary key(contactPhone), mycounter as integer mr_counter,
address.counter as integer mr_counter) in regions FRA, LON as json
collection: Error: Error found when creating the table: MR Counter
in JSON Collection may contain only alphanumeric values plus the
character "_": address.counter
```

- Ensure that all the regions mentioned in the create table statement are predefined. Declaring regions that are not available during the table creation will result in an error.
- Adding or removing the MR_COUNTERS from the JSON collection table through schema evolution is not supported.

Using the IDENTITY Column

Declare a column as IDENTITY to have Oracle NoSQL Database automatically assign values to it, where the values are generated from an associated sequence generator. The SG is the table's manager for tracking the IDENTITY column's current, next, and total number of values.

You create an IDENTITY column as part of a CREATE TABLE name DDL statement, or add an IDENTITY column to an existing table with an ALTER TABLE name DDL statement.

Syntax

```
identity_definition ::=  
GENERATED (ALWAYS | (BY DEFAULT [ON NULL])) AS IDENTITY  
["(" sequence\_generator\_attributes + ")"]
```

Semantics

An INTEGER, LONG, or NUMBER column in a table can be defined as an identity column. The system can automatically generate values for the identity column using a sequence generator. See [Sequence Generator](#) section. A value for an identity column is generated during an INSERT, UPSERT, or UPDATE statement.

An identity column can be defined either as GENERATED ALWAYS or GENERATED BY DEFAULT.

GENERATED ALWAYS

The system always generates a value for the identity column. An exception is raised if the user supplies a value for the identity column.

GENERATED BY DEFAULT

The system generates a value for the identity column only if the user does not supply a value for it. If ON NULL is specified for GENERATED BY DEFAULT, the system will generate a value when the user supplies a NULL value or the value evaluates to a NULL.

Sequence Generator Attributes

An SG has several attributes that define its behavior, such as the starting value for its IDENTITY column, or the number of values stored in cache. You can optionally define some SG attributes when you create an IDENTITY column, or use all default values. For more information about the Sequence Generator attributes, see [Sequence Generator](#).

Note

Using an IDENTITY column in any table does not force uniqueness. If your application requires unique values for every row of an IDENTITY column, you must create the column as GENERATED ALWAYS AS IDENTITY, and never permit any use of the CYCLE SG attribute.

Identity Column Characteristics

- There can be only one IDENTITY column per table.
- The IDENTITY column of a table can be part of the primary key or the shard key.
- You can add, remove, or change rows of an IDENTITY column, though certain limitations exist on such updates, depending on how you create the IDENTITY column, and whether it is a Primary Key.
- Secondary indexes can be created on an identity column.
- Dropping a table that was created with an IDENTITY column also removes the Sequence Generator.

- The set of values that may be assigned to an IDENTITY column is defined by its data type and the attributes of the sequence generator attached to it. The values are always integer numbers. Both negative and positive INTEGER are possible. If you want only positive values, then set the START WITH attribute to 1 and specify a positive INCREMENT BY attribute. When you specify CYCLE, numbers will be regenerated from the MINVALUE. In this case, if you want positive values you must also set MINVALUE to be a positive number.
- The system generates unique values for an IDENTITY column that is defined as GENERATED ALWAYS and has the sequence generator attribute NO CYCLE set. Otherwise, duplicate identity values can occur in the following scenarios:
 - The identity column is defined as GENERATED BY DEFAULT and the user supplies a value during an insert or update statement that already exists in the table for the IDENTITY column.
 - The CYCLE option is set for an IDENTITY column that is defined as GENERATED BY DEFAULT or GENERATED ALWAYS and the sequence generator reaches the end of the cycle and then recycles through the sequence generator to generate values that were generated in the previous cycle.
 - If the IDENTITY column properties are altered using the alter table statement so that during an insert or update operation the user can supply a value that already exists.
- If more than one client accesses a table with an IDENTITY column defined for unique values this way, each client is assigned contiguous value sets to its SG cache. These sets do not overlap with other client sets. For example, Client1 is assigned values 0001 – 1000, while Client2 has 1001 – 2000, and so on. Thus, as each client adds rows to the table, the IDENTITY values can run as 0001, 1001, 0002, 1002, 1003, and so on, as both clients use their own cache when adding rows. The IDENTITY column values are guaranteed to be unique, but not necessarily contiguous, because each client has its own set of cache values, and adds rows at different speeds and times.
- Sequence generator attributes can be altered using the alter table statement.
- Users require table privileges to create tables with an IDENTITY column. For a description of user privileges, see [KVStore Required Privileges](#) in the *Security Guide*.
- Holes in the sequence can occur when:
 - The application caches identity values and shuts down or crashes before using all of the cached values for inserting rows.
 - Identity values are assigned during a transaction that is rolled back.
- Creating MR tables with an IDENTITY column is not supported and the system returns an error message. For more information on MR table creation, see *Create MR Tables*.
- For example on inserting rows with an identity column see, [Inserting Rows with an IDENTITY Column](#) section.

Example 5-18 Identity Column using GENERATED ALWAYS

```
CREATE TABLE T1 (
  id INTEGER GENERATED ALWAYS AS IDENTITY
    (START WITH 2 INCREMENT BY 2 MAXVALUE 200 NO CYCLE),
  name STRING,
  PRIMARY KEY (id)
)
```

In the above example, the INTEGER column id is defined as a GENERATED ALWAYS AS IDENTITY column and is the primary key for table T. The system will start to generate values 2

through 200 incrementing by 2. So values for the id column will be 2,4,6,8,...200. Since the NO CYCLE option is defined, the system will raise an exception after the number 200 is generated saying it has reached the end of the sequence generator.

Example 5-19 Identity Column using GENERATED BY DEFAULT

```
CREATE TABLE T2 (  
    id LONG GENERATED BY DEFAULT AS IDENTITY  
        (START WITH 1 INCREMENT BY 1 CYCLE CACHE 200),  
    account_id INTEGER,  
    name STRING,  
    PRIMARY KEY (account_id)  
)
```

In the above example, the creation of a table with an identity column on id column is shown. The id column is of type LONG, is defined as GENERATED BY DEFAULT, and it is not a primary key column. This example also demonstrates how to specify a CYCLE and CACHE sequence generator attributes. The system will only generate a value during INSERT/ UPSERT/UPDATE if the user did not supply a value. It starts off generating values 1, 2, 3,... up to the maximum value of the LONG datatype, and once it exhausts all the sequence generator values, it will cycle through and re-start from the MINVALUE value of the sequence generator, which in this case, is the minimum value of the LONG datatype. The CACHE value of 200 means that every time a client uses up the values in the cache and asks for the next value, the system will give it 200 values to fill up the cache. In this example, the system will give values 1 through 200 when a client asks for a value for the first time. Another client operating on the same table may get values 201-300, so on and so forth.

Creating Tables With an IDENTITY Column

You can create an IDENTITY column when you create a table, or change an existing table to add an IDENTITY column using ALTER TABLE...ADD. In either case, choose one of the IDENTITY statements described below. This section describes creating a table with an IDENTITY column.

Here is the formal syntax for creating a table with an IDENTITY column:

```
GENERATED (ALWAYS | (BY DEFAULT [ON NULL])) AS IDENTITY  
    [sequence_options,...]
```

The optional sequence_options refer to all of the Sequence Generator attributes you can supply.

IDENTITY Column Statement	Description
GENERATED ALWAYS AS IDENTITY	The sequence generator always supplies an IDENTITY value. You cannot specify a value for the column.
GENERATED BY DEFAULT AS IDENTITY	The sequence generator supplies an IDENTITY value any time you do not supply a column value.
GENERATED BY DEFAULT ON NULL AS IDENTITY	The sequence generator supplies the next IDENTITY value if you specify a NULL column value.

Example 5-20 Create a table with a column GENERATED ALWAYS AS IDENTITY from the SQL CLI

```
CREATE TABLE IF NOT EXISTS tname1 (  
  idValue INTEGER GENERATED ALWAYS AS IDENTITY,  
  acctNumber INTEGER,  
  name STRING,  
  PRIMARY KEY (acctNumber))
```

Output:

Statement completed successfully

For this table, tname1, each time you add a row to the table, the Sequence Generator (SG) updates the idvalue from its cache. You cannot specify a value for idValue. If you do not specify any sequence generator attributes, the SG uses its default values.

Example 5-21 Create a table with a column GENERATED BY DEFAULT ON NULL AS IDENTITY

```
CREATE TABLE IF NOT EXISTS tname2 (  
  idvalue INTEGER GENERATED BY DEFAULT ON NULL AS IDENTITY,  
  acctNumber INTEGER,  
  name STRING,  
  PRIMARY KEY (acctNumber))
```

Output:

Statement completed successfully

For this table, tname2, each time you add a row, the SG inserts the next available value from its cache if no value is supplied for the idvalue column, the supplied value for the idvalue column is NULL.

Example 5-22 Create a table with a column GENERATED BY DEFAULT AS IDENTITY

```
CREATE TABLE IF NOT EXISTS tname3 (  
  idvalue INTEGER GENERATED BY DEFAULT AS IDENTITY,  
  acctNumber INTEGER,  
  name STRING,  
  PRIMARY KEY (acctNumber))
```

Output:

Statement completed successfully

For this table, tname3, each time you add a row, the SG inserts the next available value from its cache if no value is supplied for the idvalue column.

Example 5-23 create a new table, sg_atts, with several SG attributes

```
CREATE Table sg_atts (  
  id INTEGER GENERATED ALWAYS AS IDENTITY
```

```
(START WITH 2
INCREMENT BY 2
MAXVALUE 200
NO CYCLE),
name STRING,
PRIMARY KEY (id))
```

Output:

Statement completed successfully

The table `sg_atts` specifies that the integer IDENTITY field (`id`) is generated always.

SG Attribute	Description
start with 2	Start the sequence value at 2.
increment by 2	Increment the sequence value by 2 for each row.
maxvalue 200	Specifies the maximum IDENTITY value. What you specify overrides the default value maxvalue, which is the upper bound of the IDENTITY datatype in use. Once the IDENTITY column reaches this value, 200, the SG will not generate any more IDENTITY values. The maximum value has been reached and the no cycle attribute is in use.
no cycle	Do not restart from 2 or with any value at all, once the column reaches the maxvalue.

Example 5-24 create another table, `sg_some_atts`, with some SG attributes

```
CREATE Table sg_some_atts (
id LONG GENERATED BY DEFAULT AS IDENTITY
(START WITH 1
INCREMENT BY 1
CYCLE
CACHE 200),
account_id INTEGER,
name STRING,
PRIMARY KEY (account_id))
```

For the `sg_some_atts` table, specify an `id` column GENERATED BY DEFAULT AS IDENTITY, but which is not the primary key.

SG Attribute or Other Detail	Description
CYCLE	Specifying CYCLE indicates that the SG should supply IDENTITY values up to either the MAXVALUE attribute you specify, or the default MAXVALUE. When the IDENTITY reaches the MAXVALUE value, the SG restarts the values over, beginning with MINVALUE, if it is specified, or with the default MINVALUE for the data type. CYCLE is orthogonal to the CACHE attribute, which indicates only how many values to store in local cache for swift access. You can set CACHE value to closely reflect the maximum value of the datatype, but we do not recommend this, due to the client cache size.

SG Attribute or Other Detail	Description
CACHE 200	The number of values that each client stores in its cache for fast retrieval. When the IDENTITY reaches the last number in the cache, the SG gets another set of values from the server automatically.
START WITH 1	The SG generates values 1, 2, 3 and so on, until it reaches the maximum value for a LONG data type.
INCREMENT BY 1	The SG increments each new IDENTITY value for every new row.

For a full list of all sequence generator attributes, see [Sequence Generator](#).

Using the UUID data type

Overview of the UUID data type

A universally unique identifier (UUID) is a 128-bit number used to identify information in computer systems. You can create a UUID and use it to uniquely identify something. In its canonical textual representation, the 16 octets of a UUID are represented as 32 hexadecimal (base-16) digits, displayed in five groups separated by hyphens, in the form 8-4-4-4-12 for a total of 36 characters (32 hexadecimal characters and 4 hyphens). For example, a81bc81b-dead-4e5d-abff-90865d1e13b1.

In Oracle NoSQL, UUID values are represented by the UUID data type. The UUID data type is considered a subtype of the STRING data type, because UUID values are displayed in their canonical textual format and, in general, behave the same as string values in the various SQL operators and expressions. However, in order to save disk space, the UUID value is saved in a compact format on disk. If the UUID value is the primary key, the canonical 36-byte string is converted to a 19-byte string, then is saved on disk. If the UUID value is a non-primary key, the canonical 36-byte string is converted to a 16-byte array, then is saved on disk.

A table column can be declared as having UUID type in a CREATE TABLE statement. The UUID data type is best-suited in situations where you need a globally unique identifier for the records in a table that span multiple regions since identity columns are only guaranteed to be unique within a NoSQL cluster in a region.

Using the UUID data type:

Declare a column with UUID data type. UUID is a subtype of the STRING data type. This UUID column can be defined as GENERATED BY DEFAULT. The system then automatically generates a value for the UUID column if you do not supply a value for it.

Syntax:

```
uuid_definition := AS UUID [GENERATED BY DEFAULT]
```

Semantics

Declares the type of a column to be the UUID type. If the GENERATED BY DEFAULT keywords are used, the system generates a value for the UUID column automatically, if the user does not supply one.

UUID Column Characteristics :

- One table can have multiple columns defined as "STRING AS UUID". However, one table can have only one column defined as "STRING AS UUID GENERATED BY DEFAULT".

- Since the Identity column is also generated by the system, the Identity column and the UUID GENERATED BY DEFAULT columns are mutually exclusive. That means only one IDENTITY column or one "UUID GENERATED BY DEFAULT" can exist per table.
- You create a UUID column as part of a CREATE TABLE DDL statement or add a UUID column to an existing table with an ALTER TABLE DDL statement.
- You can also index UUID columns via secondary indexes.

Example 5-25 UUID Column without GENERATED BY DEFAULT

```
CREATE TABLE myTable (id STRING AS UUID, name STRING, PRIMARY KEY (id))
```

Output:

Statement completed successfully

In the above example, the id column has no "GENERATED BY DEFAULT" defined, therefore, whenever you insert a new row, you need to explicitly specify a value for the id column.

```
INSERT INTO myTable values("a81bc81b-dead-4e5d-abff-90865d1e13b1", "test1")
```

Output:

Statement completed successfully

Input format: The input string must conform to the format specified by RFC 4122. An `IllegalArgumentException` is thrown if the input string does not conform to the string representation as described at Class `UUID`.

Output format: The output is a UUID canonical format. This is 32 hexadecimal(base-16) digits, displayed in five groups separated by hyphens, in the form 8-4-4-4-12 for a total of 36 characters (32 hexadecimal characters and 4 hyphens).

The value for a UUID column can also be generated using the `random_uuid` function, which returns a randomly generated UUID, as a string of 36 characters. See Function to generate a UUID string.

Example 5-26 UUID Column using GENERATED BY DEFAULT

```
CREATE TABLE myTable (id STRING AS UUID GENERATED BY DEFAULT, name STRING,  
PRIMARY KEY (id))
```

Output:

Statement completed successfully

In the above example, the id column has "GENERATED BY DEFAULT" defined, therefore, whenever you insert a new row without specifying the value for the id column, Oracle NoSQL Database automatically generates a value for it.

```
INSERT INTO myTable VALUES(default, "test1")
```

Output:

Statement completed successfully

Table 5-2 Comparison between Identity Column and UUID column

Identity Column	UUID column
Declare a column as Identity to have Oracle NoSQL Cluster automatically assign values to it	Declare a column as UUID if you need unique values to be assigned to a NoSQL Cluster column in a multi-region system
An INTEGER, LONG, or NUMBER column in a table can be defined as an Identity column	A UUID is a subtype of the STRING data type
An Identity column can be defined either as GENERATED ALWAYS or GENERATED BY DEFAULT	A UUID column can be defined as GENERATED BY DEFAULT or you can supply the value of the string while inserting or updating data
Ideal in a single cluster architecture	The UUID data type is best suited In situations where you need a globally unique identifier for the records in a table that span multiple regions since identity columns are only guaranteed to be unique within a NoSQL cluster in a region.
Costs less storage space than a corresponding UUID column. If LONG is the primary key, it costs a maximum of 10 bytes. If LONG is a non-primary key, it costs a maximum of 8 bytes.	Costs more storage space than a corresponding Identity column. If the UUID value is the primary key, it costs 19-bytes. If the UUID value is a non-primary key, it costs 16-bytes.

Using the MR_COUNTER datatype

- [Using CRDT datatype in a multi-region table](#)
- [Create table using MR_COUNTER datatype](#)

Using CRDT datatype in a multi-region table

Overview of the MR_COUNTER data type

MR_Counter data type is a counter CRDT. CRDT stands for Conflict-free Replicated Data Type. In a multi-region setup of an Oracle NoSQL Database, a CRDT is a data type that can be replicated across servers where regions can be updated independently and it converges on a correct common state. Changes in the regions are concurrent and not synchronized with one another. In short, CRDTs provide a way for concurrent modifications to be merged across regions without user intervention. Oracle NoSQL Database currently supports the counter CRDT which is called MR_Counter. The MR_COUNTER datatype is a subtype of the INTEGER or LONG or NUMBER data type. You can also use the MR_COUNTER data type in a schema-less JSON field, which means one or more fields in a JSON document can be of MR_COUNTER data type.

Why do you need MR_Counter in a multi-region table?

In a multi-region database configuration, copies of the same data need to be stored in multiple regions. This configuration needs to deal with the fact that the data may be concurrently modified in different regions.

Take an example of a multi-region table in three different regions (where data is stored in three different Oracle NoSQL Database stores). Concurrent updates of the same data in multiple

regions, without coordination between the machines hosting the regions, can result in inconsistencies between the regions, which in the general case may not be resolvable. Restoring consistency and data integrity when there are conflicts between updates may require some or all of the updates to be entirely or partially dropped. For example, in the current configuration of a multi-region table in the Oracle NoSQL Database, if the same column (a counter) of a multi-region table is updated across two regions at the same time with different values, a conflict arises.

Currently, the conflict resolution is that the latest write overwrites the value across regions. For example, Region 1 updates column1 with a value R1, and region2 updates column1 with a value R2, and if the region2 update happens after region1, the value of the column (counter) in both the regions becomes R2. This is not what is actually desired. Rather every region should update the column (a counter) at their end and also the system internally needs to determine the sum of the column across regions.

One way to handle this conflict is making serializable/linearizable transactions (one transaction is completed and changes are synchronized in all regions and only then the next transaction happens). A significant problem of having serializable transactions is performance. This is where MR_COUNTER datatype comes in handy. With MR_COUNTER datatype, we don't need serializable transactions and the conflict resolution is taken care of. That is, MR_COUNTER datatype ensures that though data modifications can happen simultaneously on different regions, the data can always be merged into a consistent state. This merge is performed automatically by MR_COUNTER datatype, without requiring any special conflict resolution code or user intervention.

Use-case for MR_COUNTER datatype

Consider a Telecom provider providing different services and packages to its customers. One such service is a "Family Plan" option where a customer and their family share the Data Usage plan. The customer is allocated a free data usage limit for a month which your the customer's entire family collectively uses. When the total usage of customer's family reaches 90 percent of the data limit, the telecom provider sends the customer an alert. Say there are four members in customer's family plan who are spread across different physical regions. The customer needs to get an alert from the telecom provider once the total consumption of their family reaches 90 percent of the free usage. The data is replicated in different regions to cater to latency, throughput, and better performance. That means there are four regions and each has a kvstore containing the details of the customer's data usage. The usage of their family members needs to be updated in different regions and at any point in time, the total usage should be monitored and an alert should be sent if the data usage reaches the limit.

An MR_COUNTER data type is ideal in such a situation to do conflict-free tracking of the data usage across different regions. In the above example, an increment counter in every data region's data store will track the data usage in that region. The consolidated data usage for all regions can be determined by the system at any point without any user intervention. That is the total data usage at any point in time can be easily determined by the system using an MR_COUNTER datatype.

Types of MR_COUNTER Datatype

Currently, Oracle NoSQL Database supports only one type of MR_COUNTER data type. which is Positive-Negative (PN) counter.

Positive-Negative (PN) Counter

A PN counter can be incremented or decremented. Therefore, these can serve as a general-purpose counter. For example, you can use these counters to count the number of users active on a social media website at any point. When the users go offline you need to decrement the counter.

To create a multi-region table with an MR_COUNTER column, See Create multi-region table with an MR_COUNTER column section in the Administrator's Guide.

A MR_COUNTER (JSON and a non-JSON) can only be defined when the field in a schema is defined. You can do this in the following places:

- During schema definition in table creation.
- During schema definition when adding a field to the schema.

Create table using MR_COUNTER datatype

You can declare a table column of MR_Counter data type in a CREATE TABLE statement. You can do this only in a multi-region table.

Declare a column with MR_COUNTER data type. MR_COUNTER is a subtype of the INTEGER or LONG or NUMBER data type. You can also declare a field in a JSON column as MR_COUNTER data type.

Syntax:

```
mr_counter_defintion := AS MR_COUNTER
```

Semantics:

Declares the type of a column to be the MR_COUNTER type.

MR_COUNTER column Characteristics:

- MR_COUNTER data type can be used for a multi-region table only. It cannot be used in regular tables.
- One table can have multiple columns defined as "MR_COUNTER".
- You create an MR_COUNTER column as part of a CREATE TABLE DDL statement or add an MR_COUNTER column to an existing table with an ALTER TABLE DDL statement.
- You can define any field in a JSON column as an MR_COUNTER.
- The default value of an MR_COUNTER data type is always 0.
- MR_COUNTER cannot be the element of an ARRAY.

Note

MR_COUNTER cannot be a primary key or be part of a secondary index.

Example 5-27 Using MR_COUNTER data type - Create a PN counter data type in a multi-region table

```
CREATE Table myTable( name STRING,  
                      count INTEGER AS MR_COUNTER,  
                      PRIMARY KEY(name)) IN REGIONS DEN, LON
```

In the above example, you create a PN counter data type in two regions DEN and LON.

While inserting data into the table, the system inserts the default value of 0 in the following two cases.

- If you specify the "DEFAULT" keyword in the INSERT clause:
`INSERT INTO myTable VALUES ("Bob", DEFAULT);`
- If you skip the column in the INSERT clause:
`INSERT INTO myTable(name) VALUES ("Bob");`

Note

For MR_COUNTER, the count contributed by a single region cannot overflow but the MR_COUNTER value, which is the sum of counts for all regions can overflow.

For example, the above table myTable has an integer MR_COUNTER and there are two regions DEN and LON. The region DEN cannot increment the count by a total value greater than INTEGER.MAX. This is also applicable for the region LON.

However, if the region DEN increments the count by INTEGER.MAX and the region LON increments it by 2, the value of count becomes (INTEGER.MAX.+2) which has overflowed, but this is allowed.

Example 5-28 Create a multi-region table and declare fields in a JSON column as MR_COUNTER data type

```
CREATE TABLE demoJSONMR(name STRING,  
    jsonWithCounter JSON(counter as INTEGER MR_COUNTER,  
                          person.count as LONG MR_COUNTER),  
    PRIMARY KEY(name)) IN REGIONS FRA, LON
```

In the statement above, you create a multi-region table with a STRING column and a column(JSON documents). You are identifying two of the fields in the JSON document as MR_COUNTER data type. The first field is counter, which is an INTEGER MR_COUNTER data type. The second field is count within an embedded JSON document (person). The count field is of LONG MR_COUNTER data type.

Note

There is no limit on the number of MR_COUNTER data types created inside a JSON document.

Example 5-29 Insert data into multi-region table with a JSON MR_COUNTER

When inserting a row into the multi-region table with a JSON MR_COUNTER, you must supply a value 0 to the MR_COUNTER.

Note

- The system initially assigns a value of 0 to all MR_COUNTER data types even if you explicitly supply a non-zero value. This also holds good when you try to provide a value that is not an INTEGER or LONG or NUMBER.
- You can't supply the keyword DEFAULT while inserting a JSON MR_COUNTER.
- The system will return an error if you try to insert data into an MR table without supplying a value to the declared JSON MR_COUNTER field or using the keyword DEFAULT.

```
INSERT INTO demoJSONMR VALUES ("Anna",
    {
        "id" : 1,
        "counter" : 0,
        "person" : {
            "age" : 10,
            "count" : 0,
            "number" : 100
        }
    }
)
```

```
SELECT * FROM demoJSONMR
```

Output:

```
{"name":"Anna",
  "jsonWithCounter":{"id" : 1,"counter":0,
                    "person":{"age":10,"count":0,"number":100}
  }
}
```

Enabling Before-images During Table Creation

Learn how to enable before-images while creating a table.

You can create a table with an optional clause to enable the generation and persistent storage of before-images.

Syntax:

```
before_image_definition ::= ENABLE BEFORE IMAGE [ttl_definition]
```

Semantics

The ENABLE BEFORE IMAGE clause generates before-images for table rows for subsequent update and delete operations. For insert operations, before-images will be empty. In addition to enabling before-images while creating the table, you must also enable before-images in the subscription configuration through the Streams API. For more details, see Using NoSQLSubscriptionConfig in the *Streams Developer's Guide*.

If you don't include the `ENABLE BEFORE IMAGE` clause, the subscribed table can only stream after-images.

To enable before-images generation after the table is created, you must use the [ALTER TABLE Statement](#) with an `ENABLE BEFORE IMAGE` clause.

To limit the storage consumed by before-images, you can use a TTL definition to expire the before-images. The TTL value indicates that the before-images are stored on the disk from the time the before-images are generated until the TTL expires. After this duration, the before-images expire, freeing up the disk space they consumed, and will not appear in the stream.

If you don't supply a TTL definition for before-images, the generated before-images remain for 24 hours.

Note

- It is possible to set both table TTL and before-images TTL together in the `CREATE TABLE` statement. Both TTL values work independently of each other.
- It is possible that a row will expire before its before-image TTL has expired.
- Before-images are beneficial when used in conjunction with streaming. For more details, see the Before-images Streaming section.
- The before-images TTL governs the time window a before-image can be placed on the stream

Example 5-30 Enable before-images while creating airline baggage tracking application table

```
CREATE TABLE BaggageInfo (ticketNo LONG,  
fullName STRING,  
gender STRING,  
contactPhone STRING,  
confNo STRING,  
bagInfo JSON,  
PRIMARY KEY (ticketNo)  
) USING TTL 5 DAYS ENABLE BEFORE IMAGE USING TTL 48 HOURS
```

Explanation: The `CREATE TABLE` statement above creates a `BaggageInfo` table and enables before-images generation for any subsequent DML operations. The TTL value for table data is set to 5 days. The TTL value for before-images is set to 48 hours.

Sequence Generator

The sequence generator is a service that generates a sequence of integer numbers. Every `IDENTITY` column you add to a table requires its own Sequence Generator (SG). The SG is responsible for several tasks, including obtaining and supplying values to the `IDENTITY` column as necessary.

Each `IDENTITY` column requires an associated, dedicated Sequence Generator (SG). When you add an `IDENTITY` column, the system creates an SG that's runs on the client with the application. Information about all attributes for every SG is added to a system table,

SYS\$SGAttributesTable. You can see the contents of this system table using a simple query such as this:

```
SELECT * FROM SYS$SGAttributesTable
```

For other commands:

Differences in Commands	Description
SHOW TABLES	Returns a list of tables
DESCRIBE TABLE name1	Shows the schema of table name1
SELECT * FROM table_name	Shows the data rows of table_name

Syntax

```
sequence_generator_attributes ::=
  (START WITH signed_int) |
  (INCREMENT BY signed_int) |
  (MAXVALUE signed_int) | (NO MAXVALUE) |
  (MINVALUE signed_int) | (NO MINVALUE) |
  (CACHE INT) | (NO CACHE) |
  CYCLE | (NO CYCLE)
```

Semantics

Oracle NoSQL Database only supports sequence generators that are attached to identity columns. See [Using the IDENTITY Column](#). The numbers in the generated sequence depend on the attributes of the sequence generator attributes.

Following are the SG attributes that you can optionally specify when you create an IDENTITY column, or change later using the ALTER TABLE statement.

Attribute	Type	Description
START WITH	Integer	The first value in the sequence. Zero (0) is permitted as a Start With value. Default value: 1
INCREMENT BY	Integer	The next value in the sequence is generated by adding INCREMENT BY value to the current value of the sequence. The increment value can be a positive number or a negative number. Zero (0) is not permitted as an Increment By value. Specifying a negative number for Increment By decrements value from the current value of the sequence. Default value: 1

Attribute	Type	Description
MINVALUE	Integer	The lower bound of the IDENTITY values that the SG supplies. You can either specify MINVALUE or NO MINVALUE, but not both. Default value: -2^{31} , which is the minimum value of the INTEGER datatype.
NO MINVALUE	Integer	Specifies that there is no lower bound of the IDENTITY values that the SG supplies. SG uses the minimum value of the INTEGER datatype, which is -2^{31} , as the lower bound of the IDENTITY values. You can either specify MINVALUE or NO MINVALUE, but not both.
MAXVALUE	Integer	The upper bound of the IDENTITY values that the SG supplies. You can either specify MAXVALUE or NO MAXVALUE, but not both. Default value: $2^{31}-1$, which is the maximum value of the INTEGER datatype.
NO MAXVALUE	Integer	Specifies that there is no upper bound of the IDENTITY values that the SG supplies. SG uses the maximum value of the INTEGER data type, which is $2^{31}-1$, as the upper bound of the IDENTITY values. You can either specify MAXVALUE or NO MAXVALUE, but not both.

Attribute	Type	Description
CACHE	Integer	The value of this attribute specifies the count of sequence numbers that will be generated every time a request is made to the sequence generator. These requests originate at the Oracle NoSQL Database clients and are serviced by the sequence generator, which "lives" at the server. Specifically, the numbers generated in each request are sent back to the client and are cached there. Whenever a client needs to assign a value to an IDENTITY column, the next sequence number from the cache is consumed. When cache empties, a request for another batch of CACHE sequence numbers is sent to the sequence generator. If more than one client accesses a table with an IDENTITY column defined for unique values this way, each client is assigned contiguous value sets to its SG cache. These sets do not overlap with other client sets. For example, Client1 is assigned values 0001 – 1000, while Client2 has 1001 – 2000, and so on. Thus, as each client adds rows to the table, the IDENTITY values can run as 0001, 1001, 0002, 1002, 1003, and so on, as both clients use their own cache when adding rows. The IDENTITY column values are guaranteed to be unique, but not necessarily contiguous, because each client has its own set of cache values, and adds rows at different speeds and times. You can either specify CACHE or NO CACHE, but not both. Default value: 1000
NO CACHE	Integer	Specifies that local cache is not being used by the SG. You can either specify CACHE or NO CACHE, but not both.

Attribute	Type	Description
CYCLE or NO CYCLE	Boolean	Determines whether or not SG continues to generate values after reaching either the maximum or minimum value for the datatype of the IDENTITY column. After an ascending sequence reaches its maximum value, it generates its minimum value. After a descending sequence reaches its minimum, it generates its maximum value. If you specify CYCLE attribute,the SG uses the total number of values that can be generated for an IDENTITY column of a specific datatype (INTEGER, LONG, or NUMBER), unless you specify MAXVALUE to set a different limit. Once the sequence generator reaches the end of the cycle and then recycles through the sequence generator to generate values that were generated in the previous cycle. If you specify NO CYCLE attribute, Oracle NoSQL Database guarantees that each IDENTITY column value is unique, but not necessarily sequential. For example, if you set MaxValue as 10000, and multiple clients add rows to the table, each client is assigned a certain amount of values to use. Default Value: NO CYCLE

i Not
e

If you specify the CYCLE attribute, all of the existing values for the IDENTITY column

Attribute	Type	Description
		mn are used again , poten tially over writin g curre nt value s in place . Creat ing a colu mn as GEN ERA TED ALW AYS AS IDEN TITY, and using the SG NO CYCL E attrib ute is the ONL Y way to maint ain uniqu e IDEN TITY colu mn value s.

Following are internal SG attributes. You cannot specify any of these when you create or add an IDENTITY column. Each is derived from how you create the IDENTITY field. For example, one internal attribute is SGName, which is the column name you give the IDENTITY field.

Attribute	Type	Description
SGType	String	[INTERNAL EXTERNAL]. The IDENTITY column you create, or add to a table with a DDL statement. The default is INTERNAL.
SGName	String	Name of the IDENTITY field you create and with which the SG is associated.
Datatype	String	Sequence Generator datatype that you specified as part of the CREATE TABLE statement for the IDENTITY column. Each IDENTITY column can be any numeric type: INTEGER, LONG, or NUMBER.
SGAttrVersion	Long	This is an internal attribute that you cannot set. It is here for future usage.

DROP TABLE Statement

The drop table statement removes the specified table and all its associated indexes from the database.

Syntax

```
drop_table_statement ::= DROP TABLE [IF EXISTS] name_path
```

Semantics

IF EXISTS

By default, if the named table does not exist then this statement fails. If the optional IF EXISTS is specified and the table does not exist then no error is reported.

IDENTITY

When a table with an identity column is dropped, the associated sequence generator is also removed.

Example 5-31 Drop Table

```
CREATE TABLE DROPTEST (id INTEGER, name STRING, PRIMARY KEY(id))
```

```
DROP TABLE DROPTEST
```

You cannot drop a parent table if there are child tables to it. To drop a parent table, first drop all of its child tables. Otherwise, the DROP statement results in an error as shown below.

```
drop table users
```

Output:

```
Error handling command drop table users: Error: User error in query:
DROP TABLE failed for table users:
Cannot remove table users, it is still referenced by child table
```

ALTER TABLE Statement

The behavior of NoSQL Database, when a schema evolves using ALTER TABLE statement, is designed primarily for large data sets (tens to hundreds of billions of records). Simply put, big data is larger and more complex data sets, especially from new data sources. These data sets are so voluminous that traditional data processing software can't manage them. But these massive volumes of data are used to address business problems you wouldn't have been able to tackle before. Therefore, when you modify the table schema with ALTER TABLE statement, NoSQL Database does not modify every record in the table and re-write them back to disk. Instead, it uses the notion of a default value, and that value gets inserted when a reader reads data that was written with a previous version of the schema.

However, in the case of identity columns, there is no way to generate a value other than by writing a record, hence when a user alters a table and adds an identity column, any reads of that column for records that were written prior to the later table yields a null value for the identity column.

You can use the alter table command to perform the following operations.

- Add schema fields to the table schema.
- Remove schema fields from the table schema.
- Modify schema fields in the table schema.
- Add region.
- Remove region.
- Modify identity definition.
- Remove identity.
- Modify the Time-To-Live value of the table.
- Enable before-image generation for any writes to the table. For details, see *before_image_statement* in the [Semantics](#) section.

You can specify only one type of operation in a single command. For example, you cannot remove a schema field and set the TTL value together.

Note

In JSON collection tables, you can use the ALTER TABLE statement to only modify the default TTL values of a table. Any other schema alteration is not supported and an error message is returned.

Syntax

```

alter_table_statement ::=
    ALTER TABLE name_path (alter_field_statements | ttl_definition)

alter_field_statements ::=
    "(" alter_field_statement ("," alter_field_statement)* ")"

alter_field_statement ::=
    add_field_statement | drop_field_statement | modify_field_statement |
alter_regions_statement | before_image_statement

add_field_statement ::=
    ADD schema_path type\_definition
    [default\_definition | identity\_definition | uuid\_definition | mr_counter_definition]
    [comment]

drop_field_statement ::= DROP schema_path

modify_field_statement ::=
    (MODIFY schema_path identity_definition) |
    (DROP IDENTITY)

alter_regions_statement ::= add_regions_statement | drop_regions_statement

before_image_statement ::= ENABLE BEFORE IMAGE [ttl_definition]

add_regions_statement ::= ADD REGIONS region_names

drop_regions_statement ::= DROP REGIONS region_names

region_names ::= region\_name ["," region_name]*

schema_path ::= init_schema_path_step ( "." schema_path_step )*
init_schema_path_step ::= id ("[" "]" )*
    schema_path_step ::= id ("[" "]" )* | VALUES "(" ")"
ttl_definition ::= USING TTL INT_CONSTANT (HOURS | DAYS)

```

Semantics

modify_field_statement

You can use the MODIFY keyword to modify only an identity column.

add_field_statement

Adding a field does not affect the existing rows in the table. If a field is added, its default value or NULL will be used as the value of this field in existing rows that do not contain it. The field to add may be a top-level field (i.e. A table column) or it may be deeply nested inside a hierarchical table schema. As a result, the field is specified via a path. The path syntax is a subset of the one used in queries and is described in the [Path Expressions](#) section.

Note

The `mr_counter_definition` parameter declares the type of a column to be the `MR_COUNTER` datatype. This data type can be used only in a multi-region table.

drop_field_statement

Dropping a field does not affect the existing rows in the table. If a field is dropped, it will become invisible inside existing rows that do contain the field. The field to drop may be a top-level field (i.e. A table column) or it may be deeply nested inside a hierarchical table schema. As a result, the field is specified via a path. The path syntax is a subset of the one used in queries and is described in the [Path Expressions](#) section.

add_regions_statement

The `add regions` clause lets you link an existing MR Table with new regions in a multi-region Oracle NoSQL Database environment. This clause is used in expanding MR Tables to new regions. See Use Case 2: Expand a Multi-Region Table in the *Administrator's Guide*.

Note

This clause will not work with MR child tables. Instead, alter the parent table to add a new region. This will automatically add the region to all the child tables in the hierarchy.

drop_regions_statement

The `drop regions` clause lets you disconnect an existing MR Table from a participating region in a multi-region Oracle NoSQL Database environment. This clause is used in contracting MR Tables to fewer regions. See Use Case 3: Contract a Multi-Region Table *Administrator's Guide*.

before_image_statement

The `ENABLE BEFORE IMAGE` clause for a table enables the generation and persistent storage of before-images for any write operations on the table. The before-image is the table row before it gets updated or deleted by a DML operation. For insert operations, the before-image will be empty.

You can use this clause to enable before-images if your table is not enabled for before-images and your application would like to receive before-images from a subscription stream

In addition to enabling before-images on the table using the SQL statement, you must also enable before-images in the subscription configuration through the Streams API. For more details on Oracle NoSQL Database Streams API, see *Streams Developer's Guide*.

If you want to follow along with the examples, create tables and insert the data as described in [Tables used in the Examples](#).

Example 5-32 Add a field to the table schema

```
ALTER TABLE stream_acct(ADD acct_balance INTEGER DEFAULT 0)
```

Explanation: In this example, you add a new field `acct_balance` with a default value of 0 to the TV streaming application. The new field is added to rows when the rows are retrieved at any time from the table. The on-disk format is updated to include the new field when a row is written back to the table. If you supply a default value in the `ALTER TABLE` statement, the value is populated into the new field as rows are retrieved, or when rows are written and your application has not supplied a value for the new field. If a default value is not supplied, a `NULL`

value is populated instead. The field can be added as a top-level field (a table column) or it may be deeply nested inside a hierarchical record, in which case the field is specified through a path.

Example 5-33 Add a middle name to the otherNames field in the users table.

```
ALTER TABLE users (ADD otherNames[].middle STRING)
```

Explanation: The otherNames field is an array of records, where each record in the array includes the first and last name fields. In this example, you use the ALTER statement to add the middle name field to the record that is nested within the otherNames array. You are adding the middle name field to a fixed schema record data type that exists within an array. Hence, you must specify the path to the field in the ALTER TABLE statement.

Example 5-34 Remove schema fields in the table schema

```
ALTER TABLE stream_acct(DROP acct_balance)
```

Explanation: In this example, you delete the acct_balance field from the TV streaming application schema. You can drop any field in the schema other than the primary key.

If you try removing the primary key field, you get an error as shown below.

```
ALTER TABLE stream_acct(DROP acct_id)
```

Output (showing error):

```
Error handling command ALTER TABLE stream_acct(DROP acct_id): Error: at (1,
27) Cannot remove a primary key field: acct_id
```

Example 5-35 Add a region

The add regions clause lets you link an existing Multi-Region Table (MR Table) with new regions in a multi-region Oracle NoSQL Database environment. You use this clause to expand MR Tables to new regions. For more information on creating new regions, see [Create Remote Regions](#).

Associate a new region with an existing MR Table using the DDL command shown below.

```
ALTER TABLE <table_name> ADD REGIONS <region_name>
```

Explanation: Here, table_name is an MR table and region_name is an existing region. For more information, see [CREATE REGION Statement](#).

Example 5-36 Remove a region

The drop regions clause lets you disconnect an existing MR Table from a participating region in a multi-region Oracle NoSQL Database environment. You use this clause to contract MR Tables to fewer regions.

To remove an MR Table from a specific region in a multi-region Oracle NoSQL Database setup, you must run the following steps from all the other participating regions.

```
ALTER TABLE <table_name> DROP REGIONS <comma_separated_list_of_regions>
```

Explanation: Here, `table_name` is an MR Table and `comma_separated_list_of_regions` is the list of regions to be dropped.

Example 5-37 Modify the Time-To-Live value of the table

Time-to-Live (TTL) is a mechanism that allows you to set a time frame on table rows, after which the rows expire automatically and are not available anymore. By default, every table that you create has a TTL value of zero, indicating that it has no expiration time.

You can use ALTER TABLE statement to change the TTL value for any table. You can specify the TTL with a number, followed by either HOURS or DAYS.

```
ALTER TABLE stream_acct USING TTL 15 DAYS
```

Explanation: In the above statement, you add an expiry of 15 days to the new rows that get added to the TV streaming application table.

Note

Altering the TTL value for a table does not change the TTL value for existing rows in the table. Rather, it will only change the default TTL value of the rows created after altering the table. To modify the TTL of every record in a table, you must iterate through each record of the table and update its TTL value.

Example 5-38 Enable before-images for table write operations

Use the following DDL statement to enable before-images generation for any writes to a table:

```
ALTER TABLE stream_acct ENABLE BEFORE IMAGE USING TTL 10 DAYS
```

The statement above enables before-images generation with a TTL value of 10 days. The before-images for writes on the `stream_acct` table are stored on the disk for 10 days after they are generated. After this duration, the before-images expire freeing up the disk space and will not appear in the stream.

You can also enable before-images without a TTL definition as follows:

```
ALTER TABLE stream_acct ENABLE BEFORE IMAGE
```

The statement above enables before-images generation on the `stream_acct` table. The generated before-images remain for 24 hours, unless adjusted by a TTL value.

Note

- The expiration of table rows is not impacted when you enable before-images with a TTL definition.
- It is possible that a row will expire before its before-image TTL has expired.
- Before-images are beneficial when used in conjunction with streaming. For more details, see the Before-images Streaming section.
- The before-images TTL governs the time window a before-image can be placed on the stream.
- If you update a preexisting before-images TTL, then the before-images that are generated after the TTL update reflect the new expiry. The existing before-images expiry remain unaffected.

You can use the following DDL statement to disable the generation of before-images:

```
ALTER TABLE stream_acct DISABLE BEFORE IMAGE
```

The statement above disables before-images generation on the `stream_acct` table. There is no impact to after-images.

Altering an IDENTITY Column

Use the `ALTER TABLE...MODIFY` clause to change one or more attributes of a table's IDENTITY column and its Sequence Generator (SG) options.

Note

The `MODIFY` clause in an `ALTER TABLE...` statement is supported *only* on IDENTITY columns.

There are two ways to alter an IDENTITY column:

- The property of the IDENTITY column can be altered. Additionally, the sequence generator attributes associated with an IDENTITY column can be altered. The sequence generator is modified immediately with the new attributes, however, a client will see the effects of the new attributes on the sequence numbers generated on subsequent requests by the client to the sequence generator, which will happen when the cache is used up or the attributes stored at the client time out.

Note

Each client has a time-based cache to store the sequence generator attributes. The client connects to the server to refresh this cache after it expires. The default timeout is 5 mins and it can be changed by setting `sgAttrCacheTimeout` in `KVStoreConfig`.

- The IDENTITY property of an existing IDENTITY column can be dropped. The sequence generator attached to that IDENTITY column is also removed. The system will no longer generate a value for that column.

The following example shows how to alter the property of the identity column `id` from GENERATED ALWAYS to GENERATED BY DEFAULT and altering sequence generator attributes START WITH, INCREMENT BY, MAXVALUE and CACHE.

Example 5-39 To Alter the Property and Sequence Generator Attributes of an IDENTITY Column

```
CREATE Table Test_alter (  
  id INTEGER GENERATED ALWAYS AS IDENTITY  
    (START WITH 1  
     INCREMENT BY 2  
     MAXVALUE 100  
     CACHE 10  
     CYCLE),  
  name STRING, PRIMARY KEY (id)  
)
```

```
ALTER TABLE Test_alter (MODIFY id GENERATED BY DEFAULT AS IDENTITY  
(START WITH 1000  
 INCREMENT BY 3  
 MAXVALUE 5000  
 CACHE 1  
 CYCLE)  
)
```

Explanation:

Output: The IDENTITY property and the sequence generator attributes are modified.

Statement completed successfully

Example 5-40 To Drop the IDENTITY Property of an Existing IDENTITY column

```
CREATE Table Test_alter_table (  
  id INTEGER GENERATED ALWAYS AS IDENTITY(  
    START WITH 1  
    INCREMENT BY 2  
    MAXVALUE 100  
    CACHE 10  
    CYCLE),  
  name STRING,  
  PRIMARY KEY (id))
```

```
ALTER TABLE Test_alter_table (MODIFY id DROP IDENTITY)
```

Explanation: `id`

Output: The IDENTITY property and the sequence generator associated with the column `id` are dropped.

Statement completed successfully

Add or Remove a UUID column

An existing table can be altered and a new UUID column can be added. The existing records in the table will have a NULL value for the newly added UUID column. An existing UUID column can also be removed from a table.

Example 5-41 Add a UUID column to an existing table

Use ALTER TABLE to add a UUID column to an existing table.

Create a table `test_alter` without a UUID column.

```
CREATE TABLE test_alter(id INTEGER,  
                        name STRING, PRIMARY KEY(id))
```

Output:

Statement completed successfully

Use ALTER TABLE to add a UUID column to `test_alter`. You can specify the default clause, GENERATED BY DEFAULT.

```
ALTER TABLE test_alter  
            (ADD new_id STRING AS UUID GENERATED BY DEFAULT )
```

Output:

Statement completed successfully

Example 5-42 Drop a UUID column

To remove a UUID column from a table, use ALTER TABLE with a DROP id clause.

Note

You cannot drop a UUID column if it is the primary key, or if it participates in an index.

```
CREATE Table Test_alter ( name STRING ,  
                        id STRING AS UUID GENERATED BY DEFAULT,  
                        PRIMARY KEY (name))
```

Output:

Statement completed successfully

```
ALTER TABLE Test_alter (DROP id)
```

Output:

Statement completed successfully

Add or Remove an IDENTITY column

An existing table can be altered and an IDENTITY column can be added. An existing IDENTITY column can also be removed from a table.

Adding an IDENTITY Column to an Existing Table

Use ALTER TABLE to add an IDENTITY column to an existing table.

Example 5-43 Add an IDENTITY column

Create a table, test_alter, without an IDENTITY column:

```
CREATE Table test_alter
(id INTEGER,
name STRING,
PRIMARY KEY (id))
```

Output:

Statement completed successfully

Use ALTER TABLE to add an IDENTITY column to test_alter. Also specify several Sequence Generator (SG) attributes for the associated new_id IDENTITY column, but do not use the IDENTITY column as a PRIMARY KEY:

```
ALTER Table Test_alter
(ADD new_id INTEGER GENERATED ALWAYS AS IDENTITY
(START WITH 1
INCREMENT BY 2
MAXVALUE 100
CACHE 10
CYCLE))
```

Output

Statement completed successfully

Note

To add an IDENTITY column to a table, the table must be at a top level. You cannot add an IDENTITY column as the column of a deeply embedded structured datatype. Adding a column does not affect the existing rows in the table, which get populated with the new column's default value (or NULL).

Dropping an IDENTITY Column

To remove an IDENTITY column from a table, so no such field remains, use the ALTER TABLE statement.

Example 5-44 Drop an IDENTITY column

Create a table, test_alter, with an IDENTITY column:

```
CREATE Table Test_alter (  
  id INTEGER GENERATED ALWAYS AS IDENTITY(  
    START WITH 1  
    INCREMENT BY 2  
    MAXVALUE 100  
    CACHE 10  
    CYCLE),  
  name STRING,  
  PRIMARY KEY (name))
```

To drop the id column, use the ALTER TABLE statement with a DROP id clause.

```
ALTER TABLE Test_alter (DROP id)
```

Add or Remove an MR_COUNTER column

You can alter an existing multi-region table to add an MR_COUNTER column. The existing records in the table will have a value of zero for the newly added MR_COUNTER column. You can also remove an existing MR_COUNTER column from a multi-region table.

Example 5-45 Add an MR_COUNTER column to an existing multi-region table

Use ALTER TABLE to add an MR_COUNTER column to an existing multi-region table.

Create a multi-region table myTable without an MR_COUNTER column.

```
CREATE TABLE myTable ( id INTEGER,  
                        name STRING,  
                        team STRING,  
                        PRIMARY KEY (id)) IN REGIONS DEN,LON
```

Use ALTER TABLE to add an MR_COUNTER column to the table myTable. This is a PN counter.

```
ALTER TABLE myTable (ADD count INTEGER AS MR_COUNTER);
```

You can use ALTER TABLE to add a JSON MR_COUNTER field to a multi-region table.

```
ALTER TABLE demoJSONMR ADD (secondJSON  
    JSON(new_counter AS NUMBER MR_COUNTER))
```

Example 5-46 Drop an MR_COUNTER column

To remove an MR_COUNTER column from a multi-region table, use ALTER TABLE statement with a DROP id clause.

```
CREATE Table myTable( name STRING,  
    count INTEGER AS MR_COUNTER,  
    PRIMARY KEY(name)) IN REGIONS DEN,LON
```

```
ALTER TABLE myTable (DROP count)
```

You can also drop a JSON column containing MR_COUNTERS from a table using ALTER TABLE syntax.

```
ALTER TABLE demoJSONMR(DROP secondJSON)
```


6

SQL Query Management

You can use a query to retrieve data from one or more tables. A query is a statement that consists of zero or more variable declarations followed by single SELECT expression. The result of a query is always a sequence of records having the same record type.

Note

Subqueries are not supported in Oracle NoSQL Database.

Syntax

query ::= [[variable_declaration](#)] [select_expression](#)

Variable declarations and expressions will be defined later in this chapter. Before doing so, a few general concepts and related terminology must be established first.

Expressions

An **expression** represents a set of operations to be performed in order to produce a result. The various kinds of expressions supported by Oracle NoSQL Database are described later in this chapter.

Expressions are built by combining other expressions and sub-expressions through operators, function calls, or other grammatical constructs. The simplest kind of expressions are constants and references to variables or identifiers.

If you want to follow along with the examples, create the airline application table and insert data as described in [Tables used in the Examples](#).

Example 6-1 Retrieve the full name, ticket number, and bag tag number for all customer baggage shipped after the 1st of Jan 2019 in an airline application

```
SELECT fullName,  
ticketNo,  
bag.bagInfo.tagNum  
FROM BaggageInfo bag  
WHERE EXISTS bag.bagInfo[$element.bagArrivalDate >="2019-01-01T00:00:00"]
```

Explanation: This query demonstrates the usage of several expressions. In the BaggageInfo table, you store the baggage arrival date for every passenger in the bagArrivalDate field of the bagInfo array. Since the bagInfo is a JSON field, the bagArrivalDate attribute contains the date in ISO-8601 format as a string value. You use the SELECT, FROM, and WHERE clauses of the [SELECT Expression](#) to choose columns from the table.

To fetch the list of the customer baggage that was shipped after a specified date, you use the [value comparison operator](#) ">=" to check if the bagArrivalDate for each bag is greater than or equal to the given date, that is, the 1st of Jan 2019. Here, the input date is also a string in

ISO-8601 format. Using string-formatted dates in ISO-8601 format works with logical comparison operators due to the natural sort order of strings. The chronological ordering of dates happens to follow the natural sort order of strings, so in this specific case, you get the desired behavior. For more complex date operations such as date arithmetic, you must cast the ISO-8601 date string into a Timestamp data type. For more information on using the CAST operator, see [Cast Expression](#).

The [EXISTS](#) operator, which is a [comparison expression](#) returns either a true or false value based on the result of the value comparison operation. In the above query, if the bag arrival date is greater than or equal to the given date, the EXISTS operator returns a true value and the corresponding row is fetched. Else, the row is filtered out.

In this example, the bag is a table alias for the BaggageInfo table. The \$element is an implicitly-declared variable and is bound to the context row (every baggage of the customer). The variable references are a part of [primary expressions](#).

Output:

```
{ "fullName": "Fallon
Clements", "ticketNo": 1762350390409, "tagNum": "17657806255507" }
{ "fullName": "Lucinda
Beckman", "ticketNo": 1762320569757, "tagNum": "17657806240001" }
{ "fullName": "Elane Lemons", "ticketNo": 1762324912391, "tagNum": "1765780623244" }
{ "fullName": "Gerard Greene", "ticketNo": 1762341772625, "tagNum": "1765780626568" }
{ "fullName": "Kendal
Biddle", "ticketNo": 1762377974281, "tagNum": "17657806296887" }
{ "fullName": "Zulema
Martindale", "ticketNo": 1762340579411, "tagNum": "17657806288937" }
{ "fullName": "Mary Watson", "ticketNo": 1762340683564, "tagNum": "17657806299833" }
{ "fullName": "Teena Colley", "ticketNo": 1762357254392, "tagNum": "17657806255823" }
{ "fullName": "Lorenzo Phil", "ticketNo": 1762320369957, "tagNum":
[ "17657806240001", "17657806340001" ] }
{ "fullName": "Adam
Phillips", "ticketNo": 1762344493810, "tagNum": "17657806255240" }
{ "fullName": "Adelaide
Willard", "ticketNo": 1762392135540, "tagNum": "17657806224224" }
{ "fullName": "Rosalia
Triplett", "ticketNo": 1762311547917, "tagNum": "17657806215913" }
{ "fullName": "Michelle
Payne", "ticketNo": 1762330498104, "tagNum": "17657806247861" }
{ "fullName": "Zina
Christenson", "ticketNo": 1762390789239, "tagNum": "17657806228676" }
{ "fullName": "Raymond
Griffin", "ticketNo": 1762399766476, "tagNum": "17657806243578" }
{ "fullName": "Henry Jenkins", "ticketNo": 176234463813, "tagNum": "17657806216554" }
{ "fullName": "Lisbeth
Wampler", "ticketNo": 1762355854464, "tagNum": "17657806292229" }
{ "fullName": "Dierdre
Amador", "ticketNo": 1762376407826, "tagNum": "17657806240229" }
{ "fullName": "Joanne Diaz", "ticketNo": 1762383911861, "tagNum": "17657806292518" }
{ "fullName": "Omar Harvey", "ticketNo": 1762348904343, "tagNum": "17657806234185" }
{ "fullName": "Doris Martin", "ticketNo": 1762355527825, "tagNum": "17657806232501" }
```

JSON collection tables:

You can use SQL expressions to query data from JSON collection tables. The queries work in the same way as fixed schema tables. You can access the document name/value pairs in a

JSON collection table by specifying JSON path expressions. A top-level attribute in the document can be accessed using its field name as the path expression, while a nested attribute must be accessed using a path expression to the attribute. A few examples are added in the following sections to demonstrate the queries in JSON collection tables.

To follow along with the examples, create a JSON collection table for the shopping application and insert the sample data records as described in the [JSON Collection Table Example](#) section.

Sequences

A **sequence** is a set of zero or more items. All expressions operate on zero or more input sequences and produce an output sequence as their result.

A sequence is just a collection set of zero or more items (including NULLs). A sequence is not an item itself (so no nested sequences) nor is it a container: there is neither a persistent data structure nor a java class at the public API level (or internally) that represents a sequence. Expressions usually operate on sequences by iterating over their items.

Note

An array is not a sequence of items. Instead, it is a single item, albeit one that contains other items in it. So, arrays are containers of items.

Although, in general, Oracle NoSQL Database expressions work on sequences and produce sequences, many of them place restrictions on the cardinality of the sequences they accept and/or produce. For example, several expressions are scalar: they require that their input sequence(s) contain no more than one item and they never produce a sequence of more than one item. Notice that a single item is considered equivalent to a sequence containing only that single item.

Boxing and Unboxing Sequence

A sequence produced by an expression E can be converted to an array by wrapping E with an array constructor : [E]. See [Array and Map Constructors](#) section. This is called **boxing** the sequence. Conversely, there are expressions that **unbox** an array: they select all or a subset of the items contained in the array and return these items as a sequence. There is no implicit unboxing of arrays; an expression must always be applied to do the unboxing. In most cases, sequence boxing must also be done explicitly, that is, the query writer must use an array constructor. There are, however, a couple of cases where boxing is done implicitly, that is, an expression (which is not an array constructor) will convert an input sequence to an array.

Note

In standard SQL the term "expression" means "scalar expression", i.e., an expression that returns exactly one (atomic) item. The only operations that can produce more than one items (or zero items) are query blocks (either as top-level queries or subqueries) and the set operators like union, intersection, etc (in these cases, the items are tuples). In Oracle NoSQL Database too, most expressions are scalar. Like the query blocks of standard SQL, the select-form-where expression of Oracle NoSQL Database returns a sequence of items. However, to navigate and extract information from complex, hierarchical data, Oracle NoSQL Database includes path expressions as well. See [Path Expressions](#) section. Path expressions are the other main source of multi-item sequences in Oracle NoSQL Database. However, if path expressions are viewed as subqueries, the Oracle NoSQL Database model is not that different from standard SQL.

If you want to follow along with the examples, create the airline application table and insert data as described in [Tables used in the Examples](#).

Example 6-2 Fetch the passenger name and tag number for all bags whose bag arrival date is greater than 2019-03-01T13:00:00Z

```
SELECT fullname,  
bag.bagInfo[].tagNum  
FROM BaggageInfo bag  
WHERE bag.bagInfo[].bagArrivalDate >any "2019-03-01T13:00:00Z"
```

Explanation: In an airline application, each piece of baggage that is checked in by the passenger is associated with a unique tag number. In the BaggageInfo table, the tag numbers are stored in the tagnum field of the bagInfo array. If the passenger has more than one piece of luggage, the baginfo array has more than one element, and each baggage has a unique tag number.

In this query, you fetch the full name and tag numbers of all such luggage whose bag arrival date is greater than the specified value. The bagArrivalDate field is a string that holds the arrival date for each baggage in ISO-8601 format. You compare the bag arrival date of each baggage with the input date value, which is also a string in the ISO-8601 format. For passengers with additional luggage, all the associated tag numbers are listed. You use the SELECT, FROM, and WHERE clauses of the [select expression](#) to choose columns from a table.

Here, the bag is a table alias for the BaggageInfo table and can be used anywhere in the SELECT statement. The bag.bagInfo[].bagArrivalDate is a sequence of more than one item. An error is returned if you use comparison operators to compare a sequence with more than one item. Hence, you use the sequence comparison operator >any to compare the sequence with the required arrival date. For more details on sequence comparison, see [Sequence Comparison Operators](#). In this example, you can compare the string-formatted dates in ISO-8601 format due to the natural sorting order of strings without having to cast into Timestamp data types.

Output:

```
{"fullname":"Elane Lemons","tagNum":"1765780623244"}  
{"fullname":"Omar Harvey","tagNum":"17657806234185"}  
{"fullname":"Henry Jenkins","tagNum":"17657806216554"}  
{"fullname":"Kendal Biddle","tagNum":"17657806296887"}
```

```

{"fullname":"Mary Watson","tagNum":"17657806299833"}
{"fullname":"Gerard Greene","tagNum":"1765780626568"}
{"fullname":"Dierdre Amador","tagNum":"17657806240229"}
{"fullname":"Lorenzo Phil","tagNum":["17657806240001","17657806340001"]}
{"fullname":"Lucinda Beckman","tagNum":"17657806240001"}
{"fullname":"Doris Martin","tagNum":"17657806232501"}

```

Sequence Types

A sequence type specifies the type of items that may appear in a sequence, as well as an indication of the cardinality of the sequence.

Syntax

sequence_type ::= [type_definition](#) [quantifier]

quantifier := "*" | "+" | "?"

Semantics

quantifier

The quantifier is one of the following:

- * indicates a sequence of zero or more items.
- + indicates a sequence of one or more items.
- ? indicates a sequence of zero or one items.
- The absence of a quantifier indicates a sequence of exactly one item.

subtype relationship

A subtype relationship exists among sequence types as well. It is defined as follows:

- The empty sequence is a subtype of all sequence types whose quantifier is * or ?
- A sequence type SUB is a subtype of another sequence type SUP (supertype) if SUB's item type is a subtype of SUP's item type, and SUB's quantifier is a subquantifier of SUP's quantifier, where the subquantifier relationship is defined by the following matrix.

The following matrix illustrates the subquantifier relationship between the quantifiers. The column heading indicate the supertype(SUP) of the quantifier. The row heading indicate the subtype (SUB) of the quantifier.

Sup Q1 Sub Q2	one	?	+	*
one	true	false	false	false
?	true	true	false	false
+	true	false	true	false
*	true	true	true	true

For example, as per the above table, ? is a superquantifier of **one** and **?**, but is not a superquantifier of **+** and *****. Similarly, * is a superquantifier of all other quantifiers.

Note

In the following sections, when we say that an expression must have (sequence) type T, what we mean is that when the expression is evaluated, the result sequence must have type T or any subtype of T. Similarly, the usual subtype-substitution rules apply to input sequences: if an expression expects as input a sequence of type T, any subtype of T may actually be used as input.

If you want to follow along with the examples, create the airline application table and insert data as described in [Tables used in the Examples](#).

Example 6-3 Fetch the details of passengers who fly from SFO/through SFO to any other location in an airline application

```
SELECT bag.fullname,
       bag.bagInfo[].tagNum,
       bag.bagInfo[].flightLegs[].fltRouteSrc
FROM   BaggageInfo bag
WHERE  bag.bagInfo[].flightLegs[].fltRouteSrc=any "SFO"
```

Explanation: In an airline application, you can get the details of all the passengers who have traveled from a specific station. The `flightLegs` array in the `bagInfo` field contains the source and destination stations for each travel leg. The `bagInfo.flightLegs[].fltRouteSrc` is a string sequence that holds the source stations from where the passengers board the flight.

In this query, you compare the source stations with the input string, SFO, to fetch the list of passengers traveling from the SFO station. Since the comparison operators cannot operate on sequences of more than one item, you use the sequence comparison operator `=any` to compare the `bagInfo.flightLegs[].fltRouteSrc` sequence with the required station. For more details on sequence comparison, see [Sequence Comparison Operators](#).

Here, the source station can either be the flights originating from SFO, or transiting through SFO station to any other destination. The `bag` is a table alias for the `BaggageInfo` table and essentially functions as a variable. The variables are bound to the context row as a whole and can be referenced inside the hierarchically structured data.

Output:

```
{"fullname":"Henry Jenkins","tagNum":"17657806216554","fltRouteSrc":
["SFO","ORD"]}
{"fullname":"Michelle Payne","tagNum":"17657806247861","fltRouteSrc":
["SFO","IST","ATH"]}
{"fullname":"Gerard Greene","tagNum":"1765780626568","fltRouteSrc":
["SFO","IST","ATH"]}
{"fullname":"Lorenzo Phil","tagNum":
["17657806240001","17657806340001"],"fltRouteSrc":
["SFO","IST","ATH","SFO","IST","ATH"]}
{"fullname":"Lucinda Beckman","tagNum":"17657806240001","fltRouteSrc":
["SFO","IST","ATH"]}
```

Variable Declaration

Syntax

```
variable_declaration ::= DECLARE (variable_name type\_definition ";")+  
  
variable_name ::= "$" id
```

External Variables

A query may start with a variable declaration section. The variables declared here are called external variables. The value of an external variable is global and constant. The values of external variables are not known in advance when the query is formulated or compiled. Instead, the external variables must be bound to their actual values before the query is executed. This is done via programmatic APIs. See *Java Direct Driver Developer's Guide*.

The type of the item bound to an external variable must be equal to or a subtype of the variable's declared type. The use of external variables allows the same query to be compiled once and then executed multiple times, with different values for the external variables each time. All the external variables that appear in a query must be declared in the declaration section. This is because knowing the type of each external variable in advance is important for query optimization.

Note

External variables play the role of the global constant variables found in traditional programming languages (e.g. final static variables in java, or const static variables in c++).

Internal Variables

Oracle NoSQL Database allows implicit declaration of internal variables as well. Internal variables are bound to their values during the execution of the expressions that declare them.

Variables (internal and external) can be referenced in other expressions by their name. In fact, variable references themselves are expressions, and together with literals, are the starting building blocks for forming more complex expressions.

Scope

Each variable is visible (i.e., can be referenced) within a scope. The query as a whole defines the global scope, and external variables exist within this global scope. Certain expressions create sub-scopes. As a result, scopes may be nested. A variable declared in an inner scope hides another variable with the same name that is declared in an outer scope. Otherwise, within any given scope, all variable names must be unique.

Note

The names of variables are case-sensitive.

Note

The following variable names cannot be used as names for external variables: \$key, \$value, \$element, and \$pos.

Example 6-4 Variable Declaration

The following query selects the first and last names of all users whose age is greater than the value assigned to the \$age variable when the query is actually executed.

```
DECLARE $age INTEGER
```

```
SELECT firstName, lastName  
FROM Users  
WHERE age > $age
```

SELECT Expression

You can query data from the tables using the SELECT expression. Multiple clauses can be used with the SELECT expression. The clauses that can be used in the SELECT expression are given in the syntax below.

Syntax

```
select_expression ::=  
  SELECT Clause  
  from clause  
  [where clause]  
  [groupby clause]  
  [orderby clause]  
  [limit clause]  
  [offset clause]
```

Semantics

The SELECT clause and the FROM clause are mandatory.

The processing of the query starts with the FROM clause, followed by the WHERE clause (if any), followed by the GROUP BY clause (if any), followed by the ORDER BY clause (if any), followed by the SELECT clause and finishing with the OFFSET and LIMIT clauses (if any). Each clause produces a set of records, which is processed by the next clause. Each clause is described in the following sections.

If you want to follow along with the examples, create tables and insert data as described in [Tables used in the Examples](#).

FROM Clause

The FROM clause is used to retrieve rows from the referenced table(s).

Syntax

```
from_clause ::= FROM (single_from_table | nested\_tables | left\_outer\_join\_tables |  
unnest\_syntax)
```

```
single_from_table ::= aliased_table_name
```

```
aliased_table_name ::=  
    (table_name | SYSTEM_TABLE_NAME) [[AS] table_alias]
```

```
table_alias ::= [$] id
```

Semantics

As shown in the syntax, the FROM clause can either reference a single table, or include a nested table clause or a left outer join clause. It can also include an unnest syntax. For nested tables, see the [Using NESTED TABLES clause to query multiple tables in the same hierarchy](#) section. To learn more about left outer joins, see [Left Outer Join \(LOJ\)](#).

unnest_syntax

You can use `unnest_syntax` to unnest one or more arrays or maps, that is to convert the arrays or maps into a set of rows. To understand how `unnest_syntax` is used in queries, see [Unnest Arrays & Maps](#).

single_from_table

In a simple FROM clause, the table is specified by its name, which may be a composite (dot-separated) name in the case of child tables. The result of the simple FROM clause is a sequence containing the rows of the referenced table.

aliased_table_name

The table name may be followed by a table alias. Table aliases are essentially variables ranging over the rows of the specified table. If no alias is specified, one is created internally, using the name of the table as it is spelled in the query, but with dot chars replaced with '_' in the case of child tables. See [Table Hierarchies](#).

Note

Table aliases are case-sensitive, like variable names.

The other clauses of the SELECT expression operate on the rows produced by the FROM clause, processing one row at a time. The row currently being processed is called the context row. The columns of the context row can be referenced in expressions either directly by their names or by the table alias followed by a dot char and the column name. See the [Column References](#) section. If the table alias starts with a dollar sign (\$), then it actually serves as a variable declaration for a variable whose name is the alias. This variable is bound to the context row as a whole and can be referenced within sub expressions of the SELECT expression. For example, it can be passed as an argument to the `expiration_time` function to get the expiration time of the context row. See the [expiration_time function](#). In other words, a table alias like `$foo` is an expression by itself, whereas `foo` is not. Notice that if this variable has the same name as an external variable, it hides the external variable. This is because the FROM clause creates a nested scope, which exists for the rest of the SELECT expression.

Example 6-5 Select all information for all the passenger records in the airline application

```
SELECT * FROM BaggageInfo
```

Explanation: In the above query, you use the FROM clause to retrieve all the passenger data from the BaggageInfo table.

Output: One sample row:

```
"ticketNo" : 1762344493810,
"fullName" : "Adam Phillips",
"gender" : "M",
"contactPhone" : "893-324-1064",
"confNo" : "LE6J4Z",
[ {
  "id" : "79039899165297",
  "tagNum" : "17657806255240",
  "routing" : "MIA/LAX/MEL",
  "lastActionCode" : "OFFLOAD",
  "lastActionDesc" : "OFFLOAD",
  "lastSeenStation" : "MEL",
  "flightLegs" : [ {
    "flightNo" : "BM604",
    "flightDate" : "2019-02-01T01:00:00",
    "fltRouteSrc" : "MIA",
    "fltRouteDest" : "LAX",
    "estimatedArrival" : "2019-02-01T03:00:00",
    "actions" : [ {
      "actionAt" : "MIA",
      "actionCode" : "ONLOAD to LAX",
      "actionTime" : "2019-02-01T01:13:00"
    }, {
      "actionAt" : "MIA",
      "actionCode" : "BagTag Scan at MIA",
      "actionTime" : "2019-02-01T00:47:00"
    }, {
      "actionAt" : "MIA",
      "actionCode" : "Checkin at MIA",
      "actionTime" : "2019-02-01T23:38:00"
    }
  ]
}, {
  "flightNo" : "BM667",
  "flightDate" : "2019-01-31T22:13:00",
  "fltRouteSrc" : "LAX",
  "fltRouteDest" : "MEL",
  "estimatedArrival" : "2019-02-02T03:15:00",
  "actions" : [ {
    "actionAt" : "MEL",
    "actionCode" : "Offload to Carousel at MEL",
    "actionTime" : "2019-02-02T03:15:00"
  }, {
    "actionAt" : "LAX",
    "actionCode" : "ONLOAD to MEL",
    "actionTime" : "2019-02-01T07:35:00"
  }, {
```

```

        "actionAt" : "LAX",
        "actionCode" : "OFFLOAD from LAX",
        "actionTime" : "2019-02-01T07:18:00"
    } ]
} ],
"lastSeenTimeGmt" : "2019-02-02T03:13:00",
"bagArrivalDate" : "2019.02.02T03:13:00"
} ]

```

Example 6-6 Fetch the initial boarding station for all passengers from the airline application

```

SELECT DISTINCT
$bag.fullname,
$bag.contactPhone,
$flt_src as SOURCE
FROM BaggageInfo $bag,
$bag.bagInfo.flightLegs[0].fltRouteSrc $flt_src
ORDER BY $bag.fullName

```

Explanation: The `fltRouteSrc` field in the `bagInfo` array of the `BaggageInfo` table includes the details of the source stations from where the passenger boards the flight. The first element in the `flightLegs` array holds the details of the initial travel leg.

In this query, you retrieve the first element from the `flightLegs` array, that is, `$bag.bagInfo[0].flightLegs[0].fltRouteSrc` values to fetch the details of the initial stations. Notice that `$bag` is the alias for the `BaggageInfo` table and `$flt_src` is the alias for `$bag.bagInfo.flightLegs[0].fltRouteSrc` field. You can alias the field names in the `FROM` clause and use them for the rest of the `SELECT` expression.

You use the [ORDER BY clause](#) to reorder the sequence of rows in the ascending order of passenger names.

Note

In this query, you get as many entries as the total number of bags. If a passenger travels with two bags, the output displays two entries. To avoid this, you can use the `DISTINCT` keyword in the `SELECT` statement. The query then returns only one copy of each set of duplicate rows selected.

Output:

```

{"fullname":"Adam Phillips","contactPhone":"893-324-1064","SOURCE":"MIA"}
{"fullname":"Adelaide Willard","contactPhone":"421-272-8082","SOURCE":"GRU"}
{"fullname":"Dierdre Amador","contactPhone":"165-742-5715","SOURCE":"JFK"}
{"fullname":"Doris Martin","contactPhone":"289-564-3497","SOURCE":"BZN"}
{"fullname":"Elane Lemons","contactPhone":"600-918-8404","SOURCE":"MXP"}
{"fullname":"Fallon Clements","contactPhone":"849-731-1334","SOURCE":"MXP"}
{"fullname":"Gerard Greene","contactPhone":"395-837-3772","SOURCE":"SFO"}
{"fullname":"Henry Jenkins","contactPhone":"960-428-3843","SOURCE":"SFO"}
{"fullname":"Joanne Diaz","contactPhone":"334-679-5105","SOURCE":"MIA"}
{"fullname":"Kendal Biddle","contactPhone":"619-956-8760","SOURCE":"JFK"}
{"fullname":"Lisbeth Wampler","contactPhone":"796-709-9501","SOURCE":"LAX"}
{"fullname":"Lorenzo Phil","contactPhone":"364-610-4444","SOURCE":"SFO"}

```

```
{"fullname":"Lucinda Beckman","contactPhone":"364-610-4444","SOURCE":"SFO"}
{"fullname":"Mary Watson","contactPhone":"131-183-0560","SOURCE":"YYZ"}
{"fullname":"Michelle Payne","contactPhone":"575-781-6240","SOURCE":"SFO"}
{"fullname":"Omar Harvey","contactPhone":"978-191-8550","SOURCE":"MEL"}
{"fullname":"Raymond Griffin","contactPhone":"567-710-9972","SOURCE":"MSQ"}
{"fullname":"Rosalia Triplett","contactPhone":"368-769-5636","SOURCE":"JFK"}
{"fullname":"Teena Colley","contactPhone":"539-097-5220","SOURCE":"MSQ"}
{"fullname":"Zina Christenson","contactPhone":"987-210-3029","SOURCE":"MIA"}
{"fullname":"Zulema Martindale","contactPhone":"666-302-0028","SOURCE":"MIA"}
```

WHERE Clause

The WHERE clause filters the rows coming from the FROM clause, returning the rows satisfying a given condition.

Syntax

where_clause ::= WHERE [expression](#)

Semantics

For each context row, the expression in the WHERE clause is evaluated. The result of this expression must have type BOOLEAN [?](#). If the result is false, or empty, or NULL, the row is skipped; otherwise the row is passed on to the next clause.

Example 6-7 Fetch the list of Female passengers from the airline application

```
SELECT
fullname,
ticketNo
FROM Baggageinfo
WHERE gender="F"
```

Explanation: In the above query, you list the name and ticket details of the female passengers from the Baggageinfo table. The WHERE clause filters the rows based on the gender field of each record. You fetch all the records that hold the entry 'F' in the gender field.

Output:

```
{"fullname":"Adelaide Willard","ticketNo":1762392135540}
{"fullname":"Elane Lemons","ticketNo":1762324912391}
{"fullname":"Michelle Payne","ticketNo":1762330498104}
{"fullname":"Doris Martin","ticketNo":1762355527825}
{"fullname":"Rosalia Triplett","ticketNo":1762311547917}
{"fullname":"Zulema Martindale","ticketNo":1762340579411}
{"fullname":"Joanne Diaz","ticketNo":1762383911861}
{"fullname":"Kendal Biddle","ticketNo":1762377974281}
{"fullname":"Mary Watson","ticketNo":1762340683564}
```

Example 6-8 Fetch the list of passengers from the airline application whose destination station is MEL

```
SELECT
fullname,
```

```
ticketNo,  
s.bagInfo.routing[] AS ROUTING  
FROM Baggageinfo s  
WHERE regex_like(s.bagInfo.routing[], ".*MEL")
```

Explanation: In an airline application, you can fetch the list of passengers bound toward specific destination stations. The routing field in the bagInfo array holds the routing details of the passengers in the Baggageinfo table. The routing information is stored in the format Source/Transit/Destination airport codes.

In this query, you fetch the details of the passengers whose destination station is MEL. You use the `regex_like` function to achieve a pattern match to the destination airport code, MEL. Since you are only looking for a specific destination station, use the combination of the period (.) metacharacter and greedy quantifier (*) to allow zero or more occurrences of any source and transit airport code. For more details on the `regex_like` function, see [Regular Expression Conditions](#).

Output:

```
{"fullname":"Zulema Martindale","ticketNo":1762340579411,"ROUTING":"MIA/LAX/  
MEL"}  
{"fullname":"Adam Phillips","ticketNo":1762344493810,"ROUTING":"MIA/LAX/MEL"}  
{"fullname":"Zina Christenson","ticketNo":1762390789239,"ROUTING":"MIA/LAX/  
MEL"}  
{"fullname":"Joanne Diaz","ticketNo":1762383911861,"ROUTING":"MIA/LAX/MEL"}
```

Example 6-9 Fetch the list of passengers carrying more than one piece of luggage in the airline application

```
SELECT  
fullname,  
ticketNo  
FROM Baggageinfo s  
WHERE EXISTS s.baginfo[1]
```

Explanation: In this query, you list the name and ticket details of the passengers traveling with more than one piece of luggage. The bagInfo array holds the information on all the luggage owned by a passenger. If a passenger owns more than one piece of luggage, there will be more than one element in the bagInfo array. You use the `EXISTS` operator in the WHERE clause to determine whether or not a second element exists in the baginfo array.

Output:

```
{"fullname":"Lorenzo Phil","ticketNo":1762320369957}
```

Alternatively, you can use the `size()` [built-in function](#) to determine the size of the baginfo array, which is the number of bags owned by a passenger. You then use the value comparison operator '>' to check if the return value exceeds 1. You will get the same output as above. For more details on the `size()` function, see [Functions on Complex Values](#).

```
SELECT  
fullname,  
ticketNo
```

```
FROM Baggageinfo s
WHERE size(s.baginfo) > 1
```

GROUP BY Clause

The GROUP BY clause is used in a SELECT statement to collect data across multiple rows and group the result by one or more columns or expressions. The GROUP BY clause is often used with aggregate functions. Oracle NoSQL Database applies the aggregate functions to each group of rows and returns a single row for each group.

Syntax

```
groupby_clause ::= GROUP BY expression ("," expression)*
```

Semantics

Each (grouping) expression must return at most one atomic value. If a grouping expression returns an empty result on an input row, that row is skipped. Equality among grouping values is defined according to the semantics of the "=" operator, with the exception that two NULL values are considered equal. See [Value Comparison Operators](#) section. Then, for each group, a single record is constructed and returned by the GROUP BY clause. If the clause has N grouping expressions, the first N fields of the returned record store the values of the grouping expressions. The remaining M fields ($M \geq 0$) store the result of zero or more aggregate functions. In general, aggregate functions iterate over the rows of a group, evaluate an expression for each such row, and aggregate the returned values into a single value per group. Oracle NoSQL Database supports many aggregate functions as described in the [Using Aggregate Functions](#) section.

Syntactically, aggregate functions are not actually listed in the GROUP BY clause, but appear in the SELECT clause instead. In fact, aggregate functions can appear only in the SELECT or ORDER BY clauses, and they cannot be nested. Semantically, however, every aggregate function that appears in the SELECT or ORDER BY list is actually evaluated by the GROUP BY clause. If the SELECT clause contains any aggregate functions, but the SELECT expression does not contain a GROUP BY clause, the whole set of rows produced by the FROM or the WHERE clauses is considered as one group and the aggregate functions are evaluated over this single group.

The implementation of the GROUP BY clause may be index-based or generic. Index-based grouping is possible only if an index exists that sorts the rows by the values of the grouping expressions. More precisely, let $e_1, e_2, \dots, e_N$ where e_i is the i th expression (i is a number in the range 1,2,3, ...N) be the grouping expressions as they appear in the GROUP BY clause (from left to right). Then, for index-based grouping, there must exist an index (which may be the primary-key index or one of the existing secondary indexes) such that for each i in 1,2,...,N, e_i matches the definition of the i -th index field. If such an index does not exist or is not selected by the query optimizer, the GROUP BY will be generic. A generic GROUP BY uses a hash table to find rows belonging to the same group and stores all groups before returning any results to the application. The hash table is stored in the client driver memory (local hash tables, of limited size, may be used at the servers as well). As a result, a generic GROUP BY may consume a large amount of driver memory. In contrast, index-based grouping exploits the row sorting provided by the index to avoid the materialization and caching of any intermediate results. It is therefore recommended to create appropriate indexes for use in GROUP BY queries. See [Using Indexes for Query Optimization](#). Finally, notice that when you use index-based grouping, the results of a grouping SELECT expression are ordered by the grouping expressions.

Example 6-10 Group by age in the application that maintains the user data

```
SELECT
age,
count(*) AS count,
avg(income) AS income
FROM users
GROUP BY age
```

Explanation: In this query, you group users by their age using the GROUP BY clause. For each group, the query returns the associated age and the average income of the users in the group. You use the [aggregate functions](#) count to count the users in each age group and avg to calculate the average income.

Output:

```
{"age":null,"count":1,"income":75000.0}
{"age":22,"count":2,"income":50000.0}
{"age":45,"count":1,"income":75000.0}
```

Example 6-11 Display the number of bags for each passenger in the airline application

```
SELECT
bag.confNo,
count(bag.bagInfo.id) AS TOTAL_BAGS
FROM Baggageinfo bag
GROUP BY bag.confNo
```

In this query, you group the data based on the confNo using the GROUP BY clause. To fetch the number of bags per passenger, you get the count of the bagInfo.id field associated with each confNo using the count [aggregate function](#).

Output:

```
{"confNo":"FH7G1W","TOTAL_BAGS":1}
{"confNo":"PQ1M8N","TOTAL_BAGS":1}
{"confNo":"XT6K7M","TOTAL_BAGS":1}
{"confNo":"DN3I4Q","TOTAL_BAGS":1}
{"confNo":"QB100J","TOTAL_BAGS":1}
{"confNo":"TX1P7E","TOTAL_BAGS":1}
{"confNo":"CG601M","TOTAL_BAGS":1}
{"confNo":"OH2F8U","TOTAL_BAGS":1}
{"confNo":"B05G3H","TOTAL_BAGS":1}
{"confNo":"ZG8Z5N","TOTAL_BAGS":1}
{"confNo":"LE6J4Z","TOTAL_BAGS":1}
{"confNo":"XT107T","TOTAL_BAGS":1}
{"confNo":"QI3V6Q","TOTAL_BAGS":2}
{"confNo":"RL3J4Q","TOTAL_BAGS":1}
{"confNo":"HJ4J4P","TOTAL_BAGS":1}
{"confNo":"CR2C8MY","TOTAL_BAGS":1}
{"confNo":"LN0C8R","TOTAL_BAGS":1}
{"confNo":"MZ2S5R","TOTAL_BAGS":1}
{"confNo":"KN4D1L","TOTAL_BAGS":1}
{"confNo":"MC0E7R","TOTAL_BAGS":1}
```

Example 6-12 Select the total baggage originating from each airport (excluding the transit baggage) in the airline application

```
SELECT $flt_src as SOURCE,  
count(*) as COUNT  
FROM BaggageInfo $bag,  
$bag.bagInfo.flightLegs[0].fltRouteSrc $flt_src  
GROUP BY $flt_src
```

Explanation: The `fltRouteSrc` field in the `flightLegs` array of the `BaggageInfo` table includes the details of the originating station in a passenger record.

In the above query, you get the total count of baggage originating from each airport. However, you don't want to consider the airports that are part of the transit. You group the data with the flight source values of the first element of the `flightLegs` array (as the first element is the source station). You then determine the count of baggage using the [count function](#).

Output:

```
{"SOURCE": "SFO", "COUNT": 6}  
{"SOURCE": "BZN", "COUNT": 1}  
{"SOURCE": "GRU", "COUNT": 1}  
{"SOURCE": "LAX", "COUNT": 1}  
{"SOURCE": "YYZ", "COUNT": 1}  
{"SOURCE": "MEL", "COUNT": 1}  
{"SOURCE": "MIA", "COUNT": 4}  
{"SOURCE": "MSQ", "COUNT": 2}  
{"SOURCE": "MXP", "COUNT": 2}  
{"SOURCE": "JFK", "COUNT": 3}
```

JSON collection table:

The following example applies the `GROUP BY` expression on a JSON collection table. Consider a sample row from the JSON collection table created for a [shopping application](#):

```
{"contactPhone": "1617114988", "address": {"Dropbox": "Presidency  
College", "city": "Kansas City", "state": "Alabama", "zip": 95065}, "cart":  
[{"item": "A4 sheets", "priceperunit": 500, "quantity": 2}, {"item": "Mobile  
Holder", "priceperunit": 700, "quantity": 1}], "firstName": "Lorenzo", "lastName": "Ph  
il", "notify": "yes", "orders": [{"EstDelivery": "2023-11-15", "item": "AG Novels  
1", "orderID": "101200", "priceperunit": 950, "status": "Preparing to dispatch"},  
{"EstDelivery": "2023-11-01", "item": "Wallpaper", "orderID": "101200", "priceperuni  
t": 950, "status": "Transit"}]}
```

Example 6-13 Fetch the statistics of the number of orders serviced per state from the storeAcct table.

```
SELECT s.address.state, count(s.orders[]) AS TOTAL_ORDERS  
FROM storeAcct s  
WHERE size(s.orders) >= 1  
GROUP BY s.address.state
```

Explanation: In the `storeAcct` table, the shopper's address is stored in the `address` field and the purchased items are stored in the `orders` array. To fetch the number of order requests per

state, you group the shopper's data by state field using the GROUP BY expression. Count the orders that have values in them by selecting only the orders that have one or more items in the orders array.

Note

To aggregate the elements of an array, you must use the array constructor `[]` to indicate that the orders field is an array.

Output:

```
{
  "state" : "Alabama",
  "TOTAL_ORDERS" : 4
}
{
  "state" : "TX",
  "TOTAL_ORDERS" : 2
}
```

Using Aggregate Functions

You can use built in aggregate functions to find information such as a count, a sum, an average, a minimum, or a maximum.

The following functions are called SQL aggregate functions, because their semantics are similar to those in standard SQL: they work in conjunction with grouping and they aggregate values across the rows of a group. The aggregate functions can be used only in the SELECT or ORDER BY clauses, and they cannot be nested.

If you want to follow along with the examples, create tables and insert the data as described in the [Tables used in the Examples](#) topic.

The following aggregate functions are supported:

- long count(*)
- long count(any*)
- number sum(any*)
- number avg(any*)
- any_atomic min(any*)
- any_atomic max(any*)
- ARRAY(any) array_collect(DISTINCT any*)
- ARRAY(any) array_collect(any*)
- long count(DISTINCT any*)

Note

All SQL aggregate function names are case sensitive.

count(*) function

The count star function returns the number of rows in a group.

Syntax:

```
long count(*)
```

Semantics:

The count star function calculates the number of records fetched by the query.

return type: long

Example 6-14 Find the total number of passengers who have contact details in their records

```
SELECT count(*) AS COUNT_PASSENGER
FROM BaggageInfo bag
WHERE length(contactPhone) > 0
```

Explanation: In an airline baggage tracking application, you can calculate the total count of passengers who have furnished their contact details. The contactPhone field in the BaggageInfo table includes the contact details of the passengers. You use the count star function to find the number of passenger records with the contactPhone entry. There is a possibility that the contact details include NULL values and empty strings. You can use the [length](#) function in the WHERE clause to exclude such rows from being counted. If a contactPhone has an empty or a NULL value, the length function returns a NULL value. You apply a value comparison to select only the rows that yield a value greater than 0, and then determine the total count of the resultant rows using the count star function.

Output:

```
{"COUNT_PASSENGER":21}
```

Example 6-15 Find the number of checked bags that arrived on the 1st of Feb 2019

```
SELECT count(*) AS COUNT_BAGS
FROM BaggageInfo bag,
EXTRACT(DAY FROM CAST (bag.bagInfo[0].bagArrivalDate AS Timestamp(0))) $t1,
EXTRACT(MONTH FROM CAST (bag.bagInfo[0].bagArrivalDate AS Timestamp(0))) $t2,
EXTRACT(YEAR FROM CAST (bag.bagInfo[0].bagArrivalDate AS Timestamp(0))) $t3
WHERE $t3=2019 AND $t2=02 AND $t1=01
```

Explanation: In the airline baggage tracking application, you can get the total count of checked bags on a particular date. The bagArrivalDate field in the BaggageInfo table contains the arrival date of the passenger's checked bags. In the above query, you count the number of rows that have the bagArrivalDate as 1st of Feb 2019 to fetch the number of checked bags on the given date. You use several filter conditions here to extract only the date part of the timestamp. You first use the [CAST](#) operator to convert the bagArrivalDate to a timestamp and then extract the date, month, and year details from the timestamp using the [EXTRACT](#) expression. You use the value comparison to determine if the day, month, and year correspond to the required date value, that is, 01,02, and 2019 respectively. You then use the [logical operator](#) AND to select only the rows that match all three conditions and count the resultant rows using the count function.

Output:

```
{"COUNT_BAGS":1}
```

count function

The count function computes its input expression on each row in a group and counts all the non-NULL values returned by these evaluations of the input expression.

Syntax:

```
long count(any*)
```

Semantics:

any: The count function accepts any parameter type as the input argument.

return type: long

The count star function can be used when you want to count the rows of a resultant query. It does not accept any other input argument. Whereas, the count function can be used when you want to enumerate the outcome of an input expression.

Example 6-16 Find the total number of checked bags that are estimated to arrive at the LAX airport at a particular time

```
SELECT $estdate as ARRIVALDATE,  
count($flight) AS COUNT  
FROM BaggageInfo $bag,  
$bag.bagInfo.flightLegs.estimatedArrival $estdate,  
$bag.bagInfo.flightLegs.flightNo $flight,  
$bag.bagInfo.flightLegs.fltRouteDest $flt_dest  
WHERE $estdate =any "2019-02-01T11:00:00Z" AND $flt_dest =any "LAX"  
GROUP BY $estdate
```

Explanation: In an airline baggage tracking application, you can get the total count of checked bags that are estimated to arrive at a particular airport and time. For each flight leg, the `estimatedArrival` field in the `flightLegs` array of the `BaggageInfo` table contains the arrival time of the checked bags and the `fltRouteDest` field contains the destination airport code. In the above query, to determine the total number of checked bags arriving at the LAX airport at a given time, you first group the data with the estimated arrival time value using the `GROUP BY` clause. From the group, you select only the rows that have the destination airport as LAX. You then determine the bag count for the resultant rows using the count function.

Here, you can compare the string-formatted dates in ISO-8601 format due to the natural sorting order of strings without having to cast them into timestamp data types.

The `$bag.bagInfo.flightLegs.estimatedArrival` and `$bag.bagInfo.flightLegs.fltRouteDest` are sequences. Since the [comparison operator](#) `'=`' cannot operate on sequences of more than one item, the [sequence comparison operator](#) `'=any'` is used instead to compare the `estimatedArrival` and `fltRouteDest` fields.

Output:

```
{"ARRIVALDATE":"2019-02-01T11:00:00Z","COUNT":2}
```

Example 6-17 Find the number of watchers for a particular show in a TV streaming application

```
SELECT count($a.contentStreamed[$element.showName = "Bienvenu"]) AS WATCHERS
FROM stream_acct $s, $s.acct_data $a
```

Explanation: In a TV streaming application, the list of shows watched by a subscriber is stored in the `contentStreamed` array. Each element of the array corresponds to a single show. The `showName` field includes the name of each show. To fetch the number of watchers for a particular show, you use the [array-filter step expression](#) in the count function. You check whether or not the `showName` field matches the given show (in this example, `Bienvenu`) for each subscriber and count the total number of such subscribers using the count function.

Output:

```
{"WATCHERS": 4}
```

sum function

The sum function computes its input expression on each row in a group and sums up all the numeric values returned by these evaluations of the input expression. In fixed schema, the sum function returns an error if you try to sum any non-numeric field. In schema-less JSON, the sum function skips any non-numeric value. If all the values of the input expression are non-numeric, a NULL value is returned. The resulting value of the sum function has type long, double, or number, depending on the type of the input items:

- If there is at least one input item of type number, the result will be a number.
- If there is at least one item of type double or float, the result will be double, else the result will be of type long.
- If the input items are a mix of long, doubles, and numbers, the result will be of type double.
- If numeric values are not returned by the sum function's input, the result is NULL.

Syntax:

```
long sum(any*)
```

Semantics:

any: The sum function accepts any parameter type as the input argument.

return type: long

Example 6-18 Find the total screen time spent on a show by a subscriber in the TV streaming application

```
SELECT sum($content.seriesInfo[].episodes[].minWatched) AS MINSWATCHED
FROM stream_acct $s, $s.acct_data.contentStreamed[] $content
WHERE acct_id = 1 AND $content.showName = "At the Ranch"
```

Explanation: In the TV streaming application, you can calculate the total screen time spent on a show by each subscriber. Each subscriber is associated with a unique account ID, stored in the `acct_id` field of the `stream_acct` table. The information of the shows is stored in the `contentStreamed` array. The `showName` field holds the name of the show and the `minWatched` field stores the time lapsed for each episode in each season of the show. In this query, you use the sum function to add the values of the `minWatched` fields of all the episodes in all the

seasons to calculate the total screen time for the subscriber with account ID 1 and show named At the Ranch.

Output:

```
{"MINSWATCHED":225}
```

avg function

The avg (average) function computes its input expression on each row in a group and sums up as well as counts all the numeric values returned by these evaluations of the input expression. Any non-numeric values are skipped. An error message is returned if the input expression does not return any numeric values. The resulting value is the division of the sum by the count. This value has type double, or number, depending on the type of the input items.

- If there is at least one input item of type number, the result will be a number, else the result will be double.
- If numeric values are not returned by avg function's input, the result is NULL.

Syntax:

```
number avg(any*)
```

Semantics:

any: The avg function accepts any parameter type as the input argument.

return type: number

Example 6-19 Find the average screen time spent on a show by a subscriber

```
SELECT avg($content.seriesInfo[].episodes[].minWatched) AS AVERAGETIME
FROM stream_acct $s, $s.acct_data.contentStreamed[] $content
WHERE acct_id = 1 AND $content.showName = "At the Ranch"
```

Explanation: In the TV streaming application, you can calculate the average screen time spent on a show by each subscriber. Each subscriber is associated with a unique account ID, stored in the acct_id field of the stream_acct table. The information of the shows is stored in the contentStreamed array. The showName field holds the name of the series and the minWatched field stores the time lapsed for each episode in each season of the show. In this query, you use the avg function on the minWatched field to calculate the average screen time spent on the show named At the Ranch by a subscriber with account id 1. The avg function first calculates the total screen time by adding the values in the minWatched fields for the given show. It then divides the sum by the number of minWatched fields to calculate the average value.

Output:

```
{"AVERAGETIME":56.25}
```

min function

The min function returns the minimum value among all the values returned by the evaluations of the input expression on each row in a group. The input expression is evaluated as follows:

1. An error is displayed if it can be determined during the compile time that the values returned by the input expression belong to a type for which an order comparison is not

defined (for example, RECORD, MAP, BINARY, or FIXED_BINARY). Otherwise, the min value for each group is initialized to NULL.

2. Let M be the current minimum value and N be the next input value. The M and N are compared using [Value Comparison Operators](#). If M is NULL, M is set to N. Else, if N is less than M, that means N can be the minimum value and hence M is set to N, and N is set to the next input value. This is continued until all the values in the input expression are compared and a minimum value is ascertained. When the values are not comparable, the following order is used:

numeric values < timestamps < strings and enums < booleans

3. If N is a record, map, array, binary, or fixed binary value, NULL, or JSON null, it is skipped and the next input value is considered.

Syntax:

any_atomic min(any*)

Semantics:

any: The min function accepts any parameter type as the input argument.

return type: [atomic data type](#)

max function

The max function returns the maximum value in all the sequences returned by the evaluations of the input expression on each row in a group. The specific rules are the same as for the max function, except that the current max value M will be replaced by the next input value N if N is not skipped and is greater than M.

Syntax:

any_atomic max(any*)

Semantics:

any: The max function accepts any parameter type as the input argument.

return type: [atomic data type](#)

Example 6-20 For a given show, find the minimum and maximum screen time

```
SELECT min($content.seriesInfo[].episodes[].lengthMin) AS MINTIME
max($content.seriesInfo[].episodes[].lengthMin) AS MAXTIME
FROM stream_acct $s, $s.acct_data.contentStreamed[] $content
WHERE $content.showName = "At the Ranch"
```

Explanation: In the TV streaming application, you can find the minimum and maximum duration of a show. The lengthMin field in the stream_acct table stores the length of each episode for a show. In this query, you use:

- The min function on the lengthMin field to fetch the duration of the episode from the show At the Ranch that has the least screen time.
- The max function on the lengthMin field to fetch the duration of the episode from the show At the Ranch that has the most screen time.

Output:

```
{"MINTIME":45, "MAXTIME":85}
```

Example 6-21 Aggregate Function - Fetch the age and average income of users from the User data table

```
CREATE INDEX idx11 ON users (age)
```

```
SELECT  
age, count(*) AS count, avg(income) AS income  
FROM users  
GROUP BY age
```

Explanation: Consider an application that maintains the user data. See *User data application table* in the [Tables used in the Examples](#) section. The users table includes multiple records of users with the same age group. You use the aggregate functions to retrieve the count of such subscribers and their average income. The above query groups users by their age, and for each age, returns the number of users with the same age group and their average income.

Output:

```
{"age":22, "count":2, "income":50000.0}  
{"age":45, "count":1, "income":75000.0}
```

array_collect(DISTINCT any*) function

The `array_collect(DISTINCT any*)` function computes the input expression on each row of a group and collects all the resulting distinct non NULL values into an array.

Syntax:

```
ARRAY(any) array_collect(DISTINCT any*)
```

Semantics:

The `array_collect` function computes the input expression on each row of a group. The input expression may be any kind of expression, except a SELECT expression. It collects all the resulting distinct values (except NULL values) into an array. The function returns the constructed array.

Note

DISTINCT causes values to be compared for equality using the semantics of the [value comparison operator](#), with the following exceptions: strings are comparable with strings only (not enums and timestamps), enums are comparable with enums only (not strings), and timestamps are comparable with timestamps only (not strings).

Example 6-22 In the TV streaming application, return the set of distinct show ids (as an array) that have been viewed in every country.

```
SELECT acct.acct_data.country,  
array_collect(distinct acct.acct_data.contentStreamed.showId) AS shows  
FROM stream_acct acct group by acct.acct_data.country
```

Explanation: The [TV streaming application](#) streams various shows that are watched by customers across the globe. Here you want a list of distinct show ids for every country. You group the data based on the country and list the show ids for every country as an array. You eliminate duplicate show ids with the distinct operator.

Output:

```
{"country": "USA", "shows": [16, 15]}  
{"country": "France", "shows": [15]}  
{"country": "Germany", "shows": [26, 15]}
```

Example 6-23 In the TV streaming application, return the set of distinct genres (as an array) that have been viewed by customers in every country.

```
SELECT acct.acct_data.country,  
array_collect(distinct acct.acct_data.contentStreamed.genres[]) AS genres  
FROM stream_acct acct group by acct.acct_data.country
```

Explanation: The [TV streaming application](#) streams various shows that are watched by customers across the globe. Here you want a list of distinct genres of shows watched by customers in every country. You group the data based on the country and for each country, you fetch the genres of all shows into a single array using `acct.acct_data.contentStreamed.genres[]`. You eliminate duplicate genres with the distinct operator.

Note

In the above query, you use `acct.acct_data.contentStreamed.genres[]` as you want the details of all the genres to be fetched in a single array. If you omit `[]` in `acct.acct_data.contentStreamed.genres`, details of the genres will be fetched as independent arrays for each row and `distinct` will be applied only on independent arrays and not in the entire set as shown below.

```
SELECT acct.acct_data.country,  
array_collect(distinct acct.acct_data.contentStreamed.genres) AS genres  
FROM stream_acct acct group by acct.acct_data.country
```

Output:

```
{"country":"USA","genres":[["comedy","drama","danish"],  
["comedy","french"]]}  
{"country":"France","genres":[["comedy","french"]]}  
{"country":"Germany","genres":[["action","crime","spanish"],  
["comedy","french"]]}
```

In the above example, you can see that the data is grouped based on the country and for each country, the genres watched by customers are fetched as a row-level array. For example, for the country USA, you see 2 arrays of genres. The `DISTINCT` is applied at the array level. Only if the arrays are identical, then `DISTINCT` fetches only one array. Else all the arrays are fetched back as shown in the result above.

Output:

```
{"country":"USA","genres":["drama","danish","comedy","french"]}  
{"country":"France","genres":["comedy","french"]}  
{"country":"Germany","genres":["spanish","comedy","action","crime","french"]}
```

array_collect(any*) function

The `array_collect(any*)` function computes the input expression on each row of a group and collects all the resulting non NULL values into an array.

Syntax:

```
ARRAY(any) array_collect(any*)
```

Semantics:

The `array_collect` function computes the input expression on each row of a group. The input expression may be any kind of expression, except a `SELECT` expression. The function collects all the resulting values (except NULL values) into an array and returns the populated array. The `array_collect` function permits duplicate values to be inserted into the array.

Example 6-24 In the TV streaming application, return the first name and last name of customers (as an array) in USA and Germany.

```
SELECT acct.acct_data.country,  
array_collect(  
{"firstName":acct.acct_data.firstName,  
"lastName":acct.acct_data.lastName}) as user_info  
FROM stream_acct acct WHERE  
acct.acct_data.country IN ('USA','Germany')  
group by acct.acct_data.country
```

Explanation: The [TV streaming application](#) streams various shows that are watched by customers across the globe. Here you want to fetch the user information (first name and last name) of customers in USA and Germany. You group the data based on the country and filter the data only for two countries USA and Germany. For each of these two countries, you fetch the first name and last name of all the customers and populate that in a single array.

Output:

```
{"country":"USA","user_info":[{"firstName":"John","lastName":"Lewis"}]}  
{"country":"Germany","user_info":[{"firstName":"Angela","lastName":"Mercel"}]}
```

long count(DISTINCT any*) function

The count function computes the input expression on each row of a group and counts all the distinct non-NULL values returned by the input expression.

Syntax:

```
long count(DISTINCT any*)
```

Semantics:

The count function computes the input expression on each row of a group. The input expression may be any kind of expression other than a subquery. The function counts all the distinct non-NULL values returned by the input expression. The return type for the count function is long.

Note

DISTINCT causes values to be compared for equality using the semantics of the [value comparison operator](#) with the following exceptions: strings are comparable with strings only (not enums and timestamps), enums are comparable with enums only (not strings), and timestamps are comparable with timestamps only (not strings).

Example 6-25 In the TV streaming application, return the count of distinct show ids for every genre viewed by customers.

```
SELECT $genre, count(distinct $content.showId) AS show_count  
FROM stream_acct acct, acct.acct_data.contentStreamed[] as $content,  
$content.genres[] as $genre, $content.showId as $showid group by $genre
```

Explanation: The TV streaming application streams various shows that are watched by customers across the globe. Here you want a count of distinct shows in every genre viewed by

customers. The details of content streamed is fetched into a single array using `acct.acct_data.contentStreamed[]`. You group the data based on the genres and for every genre, you fetch the list of showid as an array. You eliminate duplicate show ids with the `distinct` operator. Then the count of distinct show ids for every genre is returned.

Note

In the above query, if you omit `[]` in `acct.acct_data.contentStreamed`, details of the content streamed will be fetched as independent arrays for each row and `count(distinct)` will be applied only on independent arrays and not in the entire set.

Output:

```
{"genres":"crime","show_count":1}
{"genres":"action","show_count":1}
{"genres":"comedy","show_count":2}
{"genres":"spanish","show_count":1}
{"genres":"french","show_count":1}
{"genres":"drama","show_count":1}
{"genres":"danish","show_count":1}
```

Sequence Aggregate Functions

Sequence aggregate functions simply aggregate the items in their input sequence, using the same rules as their corresponding SQL aggregate function.

For example, `seq_sum()` will skip any non-numeric items in the input sequence and it will determine the actual type of the return value (long, double, or number) the same way as the SQL `sum()`. The only exception is `seq_count()`, which contrary to the SQL `count()`, will return `NULL` if any of its input items is `NULL`. Furthermore, there are no restrictions on where sequence aggregate functions can appear (for example, they can be used in the `WHERE` and/or the `SELECT` clause).

Note

Note: An array is a sequence of one item. To aggregate the elements of an array, you must unbox the array using `[]`.

The following sequence aggregate functions are supported.

- `long seq_count(any*)`
- `number seq_sum(any*)`
- `number seq_avg(any*)`
- `any_atomic seq_min(any*)`
- `any_atomic seq_max(any*)`

Note

All sequence aggregate function names are case sensitive.

seq_count function

Returns the number of items in the input sequence. The evaluation of the input expression is similar to the [count function](#).

Syntax:

```
long seq_count(any*)
```

Semantics:

any: The seq_count function accepts any parameter type as the input argument.

return type: long

seq_sum function

Returns the sum of the numeric items in the input sequence. The evaluation of the input expression is similar to the [sum function](#).

Syntax:

```
number seq_sum(any*)
```

Semantics:

any: The seq_sum function accepts any parameter type as the input argument.

return type: number

seq_avg function

Returns the average of the numeric items in the input sequence. The evaluation of the input expression is similar to the [avg function](#).

Syntax:

```
number seq_avg(any*)
```

Semantics:

any: The seq_avg function accepts any parameter type as the input argument.

return type: number

seq_min function

Returns the minimum of the items in the input sequence. The evaluation of the input expression is similar to the [min function](#).

Syntax:

```
any_atomic seq_min(any*)
```

Semantics:

any: The seq_min function accepts any parameter type as the input argument.

return type: [atomic data type](#)

seq_max function

Returns the maximum of the items in the input sequence. The evaluation of the input expression is similar to the [max function](#).

Syntax:

```
any_atomic seq_max(any*)
```

Semantics:

any: The seq_max function accepts any parameter type as the input argument.

return type: [atomic data type](#)

Example 6-26 Display an automated message regarding the number of checked bags, travel route, and flight count to a passenger in the airline baggage tracking application

```
SELECT fullName,  
b.baginfo[0].routing,  
size(baginfo) AS BAGS,  
CASE  
    WHEN seq_count(b.bagInfo[0].flightLegs.flightNo) = 1  
    THEN "You have one flight to catch"  
    WHEN seq_count(b.bagInfo[0].flightLegs.flightNo) = 2  
    THEN "You have two flights to catch"  
    WHEN seq_count(b.bagInfo[0].flightLegs.flightNo) = 3  
    THEN "You have three flights to catch"  
    ELSE "You do not have any travel listed today"  
END AS FlightInfo  
FROM BaggageInfo b  
WHERE ticketNo = 1762320369957
```

Explanation: In the [airline baggage tracking application](#), it is helpful to display a quick look-up message regarding the flight count, number of checked bags, and routing details of an upcoming travel for a passenger. The bagInfo array holds the checked bag details of the passenger. The size of the bagInfo array determines the number of checked bags per passenger. The flightLegs array in the bagInfo includes the flight details corresponding to each travel leg. The routing field includes the airport codes of all the travel fragments. You can determine the number of flights by counting the flightNo fields in the flightLegs array. If a passenger has more than one checked bag, there will be more than one element in the bagInfo array, one for each bag. In such cases, the flightLegs array in all the elements of the bagInfo field of a passenger data will contain the same values. This is because the destination of all the checked bags for a passenger will be the same. While counting the flightNo fields, you must consider only one element of the bagInfo array to avoid duplication of results. In this query, you consider only the first element, that is, bagInfo[0]. As the flightLegs array has a flightNo field for each travel fragment, it is a sequence and you determine the count of the flightNo fields per passenger using the seq_count function.

You use the [CASE](#) statement to introduce different messages based on the flight count. For ease of use, only three transits are considered in the query.

Output:

```
{"fullName":"Lorenzo Phil","routing":"SFO/IST/ATH/  
JTR","BAGS":2,"FlightInfo":"You have three flights to catch"}
```

Example 6-27 Retrieve the step tracker details for a user from the User data application

Consider the users table with an additional field stepCount to track the steps covered by each user over a duration. The table schema is as follows:

```
CREATE TABLE users ( id INTEGER,
                      firstName STRING,
                      lastName STRING,
                      otherNames ARRAY(RECORD(first STRING, last STRING)),
                      age INTEGER,
                      income INTEGER,
                      address JSON,
                      connections ARRAY(INTEGER),
                      stepCount ARRAY(LONG),
PRIMARY KEY (id) )
```

Insert the following sample data:

```
INSERT INTO users VALUES (
    10,
    "John",
    "Smith",
    [ {"first" : "Johny", "last" : "Good"}, {"first" : "Johny2", "last" :
"Brave"}, {"first" : "Johny3", "last" : "Kind"}, {"first" : "Johny4", "last" :
"Humble"} ],
    22,
    45000,
    {
        "street" : "Pacific Ave",
        "number" : 101,
        "city" : "Santa Cruz",
        "state" : "CA",
        "zip" : 95008,
        "phones" : [
            { "area" : 408, "number" : 4538955, "kind" : "work" },
            { "area" : 831, "number" : 7533341, "kind" : "home" },
            { "area" : 831, "number" : 7533382, "kind" : "mobile" }
        ]
    },
    [ 30, 55, 43 ],
    [ 2000, 1500, 2700, 3000, 1000, 4000, 6000]
)
```

Based on the aggregation of this data for a given duration, a user can chart out a fitness regime.

```
SELECT id,
firstName,
seq_count(u.stepCount[]) AS DAYS,
seq_sum(u.stepCount[]) AS TOTAL_STEPS,
seq_avg(u.stepCount[]) AS AVERAGE_STEPS,
seq_min(u.stepCount[]) AS LOWEST,
seq_max(u.stepCount[]) AS HIGHEST
```

```
FROM users u  
WHERE id=10
```

Explanation: The stepCount field in the users table is an array of type long. Each element in the array represents the number of steps covered by a user per day. You can use the sequence aggregate functions to fetch the details of total steps, average steps, lowest, and highest counts of steps covered by the user over a period.

Output:

```
{"id":10,"firstName":"John","DAYS":7,"TOTAL_STEPS":20200,"AVERAGE_STEPS":2885.  
714285714286,"LOWEST":1000,"HIGHEST":6000}
```

ORDER BY Clause

The ORDER BY clause reorders the sequence of rows it receives as input. The relative order between any two input rows is determined by evaluating, for each row, the expressions listed in the ORDER BY clause and comparing the resulting values, taking into account the sort_spec associated with each ORDER BY expression.

Syntax

```
orderby_clause ::= ORDER BY  
    expression sort_spec  
    ("," expression sort_spec)*  
  
sort_spec ::= [ASC|DESC] [NULLS (FIRST|LAST)]
```

Semantics

Each ordering expression must return at most one atomic value. If an ordering expression returns an empty sequence, the special value EMPTY is used as the returned value. If the SELECT expression includes GROUP BY as well, then the expressions in the ORDER BY must be the grouping expressions (in the GROUP BY clause, if any), or aggregate functions, or expressions that are built on top of grouping expression and/or aggregate functions.

sort_spec : A sort_spec specifies the "direction" of the sort (ascending or descending) and how to compare the special values NULL, JNULL, and EMPTY with the non-special values.

- If NULLS LAST is specified, the special values will appear after all the non-special values.
- If NULLS FIRST is specified, the special values will appear before all the non-special values.

The relative ordering among the 3 special values themselves is fixed:

- if the direction is ASC, the ordering is EMPTY < JNULL < NULL;
- otherwise the ordering is reversed.

Notice that in the grammar, sort_specs are optional.

- If no sort_spec is given, the default is ASC order and NULLS LAST.
- If only the sort order is specified, then NULLS LAST is used if the order is ASC, otherwise NULLS FIRST.
- If the sort order is not specified, ASC is used.

Taking into account the above rules, the relative order between any two input rows is determined as follows. Let N be the number of order-by expressions and let $V_{i1}, V_{i2}, \dots, V_{iN}$ be the atomic values (including EMPTY) returned by evaluating these expressions, from left to right, on a row R_i . Two rows R_i, R_j are considered equal if V_{ik} is equal to V_{jk} for each k in 1, 2, ..., N . In this context, NULLs are considered to be equal only to themselves. Otherwise, R_i is considered less than R_j if there is a pair V_{im}, V_{jm} such that:

- m is 1, or V_{ik} is equal to V_{jk} for each k in 1, 2, ..., $(m-1)$, and
- V_{im} is not equal to V_{jm} , and
- the m -th sort_spec specifies ascending order and V_{im} is less than V_{jm} , or
- the m -th sort_spec specifies descending order and V_{im} is greater than V_{jm}

In the above rules, comparison of any two values V_{ik} and V_{jk} , when neither of them is special and they are comparable to each other, is done according to the rules of the value-comparison operators defined in the [Value Comparison Operators](#) section.

If V_{ik} and V_{jk} do not have comparable types (which, for example, - can arise when sorting by json fields), the following rule applies:

- If the direction is ASC, the ordering is numeric items < timestamps < strings and enums < booleans.
- Otherwise the ordering is reversed.

As with grouping, sorting can be index-based or generic. Index-based sorting is possible only if there is an index that sorts the rows in the desired order. More precisely, let $e_1, e_2, \dots, e_N$ by the ORDER BY expressions as they appear in the ORDER BY clause (from left to right). Then, there must exist an index (which may be the primary-key index or one of the existing secondary indexes) such that for each i in 1,2,...,N, e_i matches the definition of the i -th index field. Furthermore, all the sort_specs must specify the same ordering direction and for each sort_spec, the desired ordering with respect to the special values must match the way these values are sorted by the index. In the current implementation, the special values are always sorted last in an index. So, if the sort order is ASC, all sort_specs must specify NULL LAST, and if the sort order is DESC, all sort_specs must specify NULLS FIRST.

Note

If no appropriate index exists or is not selected by the query optimizer, the sorting will be generic. This implies that all query results must be fetched into the driver memory and cached there before they can be sorted. So, as with grouping, generic sorting can consume a lot of driver memory, and is therefore best avoided.

For both generic ORDER BY and GROUP BY, applications can programmatically specify how much memory such operations are allowed to consume at the client driver. We have specific methods for this functionality in each of the available language drivers as given below.

Table 6-1 APIS for Memory Consumption

Language Driver	Get maximum memory consumption	Set maximum memory consumption
Java	<code>getMaxMemoryConsumption()</code>	<code>setMaxMemoryConsumption(long v)</code>
Python	<code>get_max_memory_consumption()</code>	<code>set_max_memory_consumption(memory_consumption)</code>

Table 6-1 (Cont.) APIS for Memory Consumption

Language Driver	Get maximum memory consumption	Set maximum memory consumption
Node.js	maxMemoryMB	maxMemoryMB
Go	GetMaxMemoryConsumption()	MaxMemoryConsumption

Example 6-28 ORDER BY Clause

This example selects the id and the last name for users whose age is greater than 30, returning the results sorted by id.

```
SELECT id, lastName
FROM users
WHERE age > 30
ORDER BY id
```

Example 6-29 ORDER BY Clause

This example selects the id and the last name for users whose age is greater than 30, returning the results sorted by age. An index-based sorting is performed if there exists a secondary index on the age column.

```
SELECT id, lastName
FROM users
WHERE age > 30
ORDER BY age
```

Example 6-30 ORDER BY Clause

The following example returns all the rows sorted by the first name.

```
SELECT id, firstName, lastName
FROM users
ORDER BY firstName
```

Output

```
+----+-----+-----+
| id | firstName | lastName |
+----+-----+-----+
| 10 | John      | Smith    |
| 20 | Mary      | Ann      |
| 30 | Peter     | Paul     |
+----+-----+-----+
3 rows returned
```

Example 6-31 ORDER BY Clause

The following example returns the firstName, lastName and income sorted by the income from highest to lowest.

```
SELECT firstName, lastName, income
   from users
  ORDER BY income DESC
```

Output

```
+-----+-----+-----+
| firstName | lastName | income |
+-----+-----+-----+
| Mary      | Ann      | 90000  |
| Peter     | Paul     | 53000  |
| John      | Smith    | 45000  |
+-----+-----+-----+
3 rows returned
```

Example 6-32 ORDER BY Clause

The following example groups the data by age and returns the number of users having that age and their average income ordered by their average income.

```
SELECT
  age, count(*), avg(income)
  FROM users
 GROUP BY age
  ORDER BY avg(income)
```

Output

```
+-----+-----+-----+
| age | Column_2 | Column_3 |
+-----+-----+-----+
| 22  |          1 | 45000.0  |
| 25  |          1 | 53000.0  |
| 43  |          1 | 90000.0  |
+-----+-----+-----+
3 rows returned
```

Example 6-33 ORDER BY Clause

In the following example, Query 1 returns the state and income sorted by income. However, if we want to group Query 1 by state, then we can use the GROUP BY clause. However, when a SELECT expression includes grouping, expressions in the SELECT and ORDER BY clauses must reference grouping expressions, aggregate functions or external variable only. So, to get the desired result, we need to rewrite Query 1 as given in Query 2.

Query 1:

```
SELECT
  u.address.state, u.income
```

```
FROM users u
ORDER BY u.income
```

Output:

```
+-----+-----+
|      state      | income |
+-----+-----+
| NV              |  45000 |
+-----+-----+
| CA              |  53000 |
+-----+-----+
| CA              |  90000 |
+-----+-----+
```

3 rows returned

Query 2:

```
SELECT
  u.address.state, max(u.income)
FROM users u
GROUP BY u.address.state
ORDER BY max(u.income)
```

Output:

```
+-----+-----+
|      state      | Column_2 |
+-----+-----+
| NV              |    45000 |
+-----+-----+
| CA              |    90000 |
+-----+-----+
```

2 rows returned

Example 6-34 ORDER BY Clause

In the following example, the Query 1 returns the income and state of all the rows in the users table. The Query 2 gets the average income for each state.

Query 1:

```
SELECT
  u.address.state, u.income
FROM users u
```

Output:

```
+-----+-----+
|      state      | income |
+-----+-----+
| CA              |  53000 |
+-----+-----+
```

NV	45000
CA	90000

3 rows returned

Query 2:

```
SELECT
    u.address.state, avg(u.income)
  FROM users u
 GROUP BY u.address.state
 ORDER BY avg(u.income)
```

Output:

state	Column_2
NV	45000.0
CA	71500.0

2 rows returned

SELECT Clause

The SELECT clause transforms each input row to a new record that will appear in the query result. The SELECT clause comes in two forms: "select star" form and "projection" form.

select star form

In select star form the SELECT clause contains a single star symbol (*). In this the SELECT clause is a no-op; it simply returns its input sequence of rows.

projection form

In the projection form the SELECT clause contains a list of expressions, where each expression is optionally associated with a name. In this the listed expressions and their associated names are referred as field expressions and field names respectively.

Syntax

```
select_clause ::= SELECT [DISTINCT] select_list
```

```
select_list ::= [hints]
              (STAR | (expression AS id ("," expression AS id)*))
```

Semantics

In projection form, the SELECT clause creates a new record for each input row. In this the record constructed by the SELECT clause has one field for each field expression and the fields are arranged in the same order as the field expressions. For each field, its value is the value computed by the corresponding field expression and its name is the name specified by the AS keyword, or if no field name is provided explicitly (via the AS keyword), one is generated internally during query compilation. To create valid records, the field names must be unique.

Furthermore, each field value must be exactly one item. To achieve this, the following two implicit conversions are employed:

1. If the result of a field expression is empty, NULL is used as the value of the corresponding field in the created record.
2. If the compiler determines that a field expression may return more than one item, it wraps the field expression with a conditional array constructor. See the [Array and Map Constructors](#) section. During runtime, an array will be constructed only if the field expression does actually return more than one item; if so, the returned items will be inserted into the constructed array, which will then be used as the value of the corresponding field in the created record.

The above semantics imply that all records generated by a SELECT clause have the same number of fields and the same field names. As a result, a record type can be created during compilation time that includes all the records in the result set of a query. This record type is the type associated with each created record, and is available programmatically to the application.

The SELECT clause can contain an optional DISTINCT keyword. If the SELECT clause contains the DISTINCT keyword, then the database will return only one copy of each set of duplicate rows selected. Duplicate rows are those with matching values in the SELECT list. The query uses the combination of values in all specified columns in the SELECT list to evaluate the uniqueness. See the [Example 6-40](#) example. Equality between values is checked using the semantics of the "=" operator. See the [Value Comparison Operators](#) section.

If the SELECT expression is a grouping one, then the expressions in the SELECT list must be the grouping expressions (in the GROUP BY clause, if any), or aggregate functions, or expressions that are built on top of grouping expression and/or aggregate functions.

Note

The SELECT clause may also contain one or more hints, that help the query processor choose an index to use for the query. See the [Choosing the Best Applicable Index](#) section.

Example 6-35 SELECT Clause

```
SELECT * FROM users
```

Example 6-36 SELECT Clause

Select the id and the last name for users whose age is greater than 30. We show 4 different ways of writing this query, illustrating the different ways that the top-level columns of a table may be accessed.

```
SELECT id, lastName FROM users WHERE age > 30
```

```
SELECT users.id, lastName FROM users WHERE users.age > 30
```

```
SELECT $u.id, lastName FROM users $u WHERE $u.age > 30
```

```
SELECT u.id, lastName FROM users u WHERE users.age > 30
```

Example 6-37 SELECT Clause

Select the id and the last name for users whose age is greater than 30, returning the results sorted by id. Sorting is possible in this case because id is the primary key of the users table.

```
SELECT id, lastName FROM users WHERE age > 30 ORDER BY id
```

Example 6-38 SELECT Clause

Select the list of distinct age of the users.

```
SELECT DISTINCT age FROM users
```

Output

```
{"age":25}
{"age":43}
{"age":22}
```

Example 6-39 SELECT Clause

Select the list of othernames of the users. Notice that the output of the SELECT command is compared and any duplicates are removed from the final output.

```
SELECT otherNames FROM Users
```

Output

```
{"otherNames":null}
{"otherNames":null}
{"otherNames":[{"first":"Johny","last":"BeGood"}]}
```

```
SELECT DISTINCT otherNames FROM Users
```

Output

```
{"otherNames":null}
{"otherNames":[{"first":"Johny","last":"BeGood"}]}
```

Example 6-40 SELECT Clause

Select the list of firstname and othernames of the users. Notice that the query uses the combination of values in all specified columns in the SELECT list to evaluate the uniqueness.

```
SELECT firstName, otherNames FROM Users
```

Output

```
{"firstName":"Peter","otherNames":null}  
{"firstName":"Mary","otherNames":null}  
{"firstName":"John","otherNames":[{"first":"Johnny","last":"BeGood"}]}
```

```
SELECT DISTINCT firstName, otherNames FROM Users
```

Output

```
{"firstName":"Peter","otherNames":null}  
{"firstName":"Mary","otherNames":null}  
{"firstName":"John","otherNames":[{"first":"Johnny","last":"BeGood"}]}
```

LIMIT Clause

The LIMIT clause is used to specify the maximum number M of results to return to the application. M is computed by an expression that may be a single integer literal, or a single external variable, or any expression which is built from literals and external variables and returns a single non-negative integer.

Syntax

```
limit_clause ::= LIMIT add\_expression
```

Semantics

Although it's possible to use limit without an order-by clause, it does not make much sense to do so. This is because without an order-by, results are returned in a random order, so the subset of results returned will be different each time the query is run.

Example 6-41 LIMIT Clause

```
SELECT * FROM users  
WHERE age > 30  
ORDER BY age  
LIMIT 5
```

OFFSET Clause

The OFFSET clause is used to specify a number N of initial query results that should be skipped (not returned to the application). N is computed by an expression that may be a single integer literal, or a single external variable, or any expression which is built from literals and external variables and returns a single non-negative integer.

Syntax

```
offset_clause ::= OFFSET add\_expression
```

Semantics

Although it's possible to use offset without an order-by clause, it does not make much sense to do so. This is because without an order-by, results are returned in a random order, so the subset of results skipped will be different each time the query is run.

Example 6-42 OFFSET Clause

```
SELECT * FROM users
WHERE age > 30
ORDER BY age
OFFSET 10
```

Path Expressions

Syntax

path_expression ::= [primary_expression](#) (map_step | array_step)*

map_step ::= "." ([map_filter_step](#) | [map_field_step](#))

array_step ::= [array_filter_step](#) | [array_slice_step](#)

Semantics

Path expressions are used to navigate inside hierarchically structured data. As shown in the syntax, a path expression has an input expression (which is one of the primary expressions described in the [Primary Expressions](#) section, followed by one or more steps. The input expression may return any sequence of items. Each step is actually an expression by itself; It takes as input a sequence of items and produces zero or more items, which serve as the input to the next step, if any. Each step creates a nested scope, which covers just the step itself.

All steps iterate over their input sequence, producing zero or more items for each input item. If the input sequence is empty, the result of the step is also empty. Otherwise, the overall result of the step is the concatenation of the results produced for each input item. The input item that a step is currently operating on is called the context item, and it is available within the step expression via an implicitly-declared variable, whose name is a single dollar sign (\$). This context-item variable exists in the scope created by the step expression.

There are several kinds of steps. For all of them, if the context item is NULL, it is just added into the output sequence with no further processing. Otherwise, the following subsections describe the operation performed by each kind of step on each non-NULL context item.

Field Step Expressions

Syntax

map_field_step ::=
[id](#) | [string](#) | [variable_reference](#) |
[parenthesized_expression](#) | [function_call](#)

Semantics

The main use of a field step is to select the value of a field from a record or map. The field to select is specified by its field name, which is either given explicitly as an identifier, or is computed by a name expression. The name expression, must have type `STRING?`.

A field step processes each context item as follows:

- If the context item is an atomic item, it is skipped (the result is empty).
- The name expression is computed. The name expression may reference the context item via the `$` variable. If the name expression returns the empty sequence or `NULL`, the context item is skipped. Otherwise, let `K` be the result of the name expression (if an identifier is used instead of a name expression, `K` is the string with the same characters as the identifier).
- If the context item is a record, then if that record contains a field whose name is equal to `K`, the value of that field is returned, otherwise, an error is raised.
- If the context item is a map, then if that map contains a field whose name is equal to `K`, the value of that field is returned, otherwise, an empty result is returned.
- If the content item is an array, the field step is applied recursively to each element of the array (with the context item being set to the current array element).

Example 6-43 Field Step Expression

Select the id and the city of all users.

```
SELECT id, u.address.city
FROM users u
```

Notice that if the input to a path expressions is a table column, a table alias must be used together with the column name. Otherwise, as explained in the [Variable References](#) section, an expression like `address.city` would be interpreted as a reference to the city column of a table called `address`, which is of course not correct.

Recall that `address` is a column of type `JSON`. For most (if not all) users, its value will be a json document, i.e., a map containing other json values. If it is a document and it has a field called `city`, its value will be returned. For address documents with no `city` field, the path expression `u.address.city` returns the empty sequence, which gets converted to `NULL` by the `SELECT` clause. The same is true for addresses that are atomic values (e.g. flat strings). Finally, a user may have many addresses stored as an array in the `address` column. For such a user, all of his/her cities will be returned inside an array.

The record items constructed and returned by the above query will all have type `RECORD(id INTEGER, city JSON)`. The `city` field of this record type has type `JSON`, because the `address` column has type `JSON` and as a result, any nested field in an address can have any valid `JSON` value. However, each actual record value in the result will have a `city` field whose field value has a more specific type (most likely `STRING`).

Note

The query processor could be constructing on-the-fly a precise `RECORD` type for each individual record constructed by the query, but it does not do so for performance reasons. Instead it constructs a common type for all returned record items.

Example 6-44 Field Step Expression

Select the id and amount spent on books for all users who live in California.

```
SELECT id, u.expenses.books
FROM users u
WHERE u.address.state = "CA"
```

In this case, "expenses" is a "typed" map: all of its values have INTEGER as their type. As a result, the record items constructed and returned by the above query will all have type RECORD(id INTEGER, books INTEGER).

Example 6-45 Field Step Expression

For each user, select their id and a field from his/her address. The field to select is specified via an external variable.

```
DECLARE $fieldName STRING

SELECT u.id, u.address.$fieldName
FROM users u
```

Example 6-46 Field Step Expression

For each user select all their last names. In this query the otherName column is an array, and the .last step is applied to each element of the array.

```
SELECT lastName, u.otherNames.last
FROM users u
```

Example 6-47 Field Step Expression

For each user select their id and all of their phone numbers (without the area code). This query will work as expected independently of whether phones is an array of phone objects or a single such phone object. However, if phones is, for example, a single integer or a json object without a number field, the path expression will return the empty sequence, which will be converted to NULL by the SELECT clause.

```
SELECT id, u.address.phones.number
FROM users u
```

Example 6-48 Field Step Expression

For each state, find how much people in that state spent on books.

```
SELECT u.address.state, sum(u.expenses.books)
FROM users u
GROUP BY u.address.state
```

For the above query to work, an index must exist whose first field is u.address.state.

Map-Filter Step Expressions

Syntax

```
map_filter_step ::= (KEYS | VALUES) "(" [expression] ")"
```

Semantics

Like field steps, map-filter steps are meant to be used primarily with records and maps. Map-filter steps select either the field names (keys) or the field values of the map/record fields that satisfy a given condition (specified as a predicate expression inside parentheses). If the predicate expression is missing, it is assumed to be the constant true (in which case all the field names or all of the field values will be returned).

A map filter step processes each context item as follows:

- If the context item is an atomic item, it is skipped (the result is empty).
- If the context item is a record or map, the step iterates over its fields. For each field, the predicate expression is computed. In addition to the context-item variable (\$), the predicate expression may reference the following two implicitly-declared variables: \$key is bound to the name of the context field, i.e., the current field in \$, and \$value is bound to the value of the context field. The predicate expression must be [BOOLEAN?](#). A NULL or an empty result from the predicate expression is treated as a false value. If the predicate result is true, the context field is selected and either its name or its value is returned; otherwise the context field is skipped.
- If the context item is an array, the map-filter step is applied recursively to each element of the array (with the context item being set to the current array element).

Example 6-49 Map-Filter Step Expressions

For each user select their id and the expense categories in which the user spent more than \$1000.

```
SELECT id, u.expenses.keys($value > 1000)
FROM users u
```

Example 6-50 Map-Filter Step Expressions

For each user select their id and the expense categories in which they spent more than they spent on clothes. In this query, the context-item variable (\$) appearing in the filter step expression [\$value > \$.clothes] refers to the context item of that filter step, i.e., to an expenses map as a whole.

```
SELECT id, u.expenses.keys($value > $.clothes)
FROM users u
```

Example 6-51 Map-Filter Step Expressions

For each user select their id, the sum of their expenses in all categories except housing, and the maximum of these expenses.

```
SELECT id,
seq_sum(u.expenses.values($key != housing)) AS sum,
```

```
seq_max(u.expenses.values($key != housing)) AS max  
FROM users u
```

Example 6-52 Map-Filter Step Expressions

Notice that field steps are actually a special case of map-filter steps. For example the query

```
SELECT id, u.address.city  
FROM users u
```

is equivalent to

```
SELECT id, u.address.values($key = "city")  
FROM users u
```

However, the field step version is the preferred one, for performance reasons.

Array-Filter Step Expressions

Syntax

```
array_filter_step ::= "[" expression "]"
```

Semantics

An array filter is similar to a map filter, but it is meant to be used primarily for arrays. An array filter step selects elements of arrays by computing a predicate expression for each element and selecting or rejecting the element depending on the predicate result. The result of the filter step is a sequence containing all selected items. If the predicate expression is missing, it is assumed to be the constant true (in which case all the array elements will be returned).

An array filter step processes each context item as follows:

- If the context item is not an array, an array is created and the context item is added to that array. Then the array filter is applied to this single-item array as described below.
- If the context item is an array, the step iterates over the array elements and computes the predicate expression on each element. In addition to the context-item variable (\$), the predicate expression may reference the following two implicitly-declared variables: \$element is bound to the context element, i.e., the current element in \$, and \$pos is bound to the position of the context element within the array (positions are counted starting with 0). The predicate expression must return a boolean item, or a numeric item, or the empty sequence, or NULL. A NULL or an empty result from the predicate expression is treated as a false value. If the predicate result is true/false, the context element is selected/skipped, respectively. If the predicate result is a number P, the context element is selected only if the condition \$pos = P is true. Notice that this implies that if P is negative or greater or equal to the array size, the context element is skipped.

Example 6-53 Array-Filter Step Expression

For each user, select their last name and his/her phone numbers with area code 650.

```
SELECT lastName,  
[ u.address.phones[$element.area = 650].number ]
```

```
AS phoneNumbers
FROM users u
```

Notice the the path expression in the select clause is enclosed in square brackets, which is the syntax used for arrayconstructor expressions as described in the [Array and Map Constructors](#) section. The use of the explicit array constructor guarantees that the records in the result set will always have an array as their second field. Otherwise, the result records would contain an array for users with more than one phones, but a single integer for users with just one phone. Notice also that for users with just one phone, the phones field in address may not be an array (containing a single phone object), but just a single phone object. If such a single phone object has area code 650, its number will be selected, as expected.

Example 6-54 Array-Filter Step Expression

For each user, select their last name and phone numbers having the same area code as the first phone number of that user.

```
SELECT lastName,
[ u.address.phones[$element.area = ${0}.area].number ]
FROM users u
```

Example 6-55 Array-Filter Step Expression

Among the 10 strongest connections of each user, select the ones with id > 100. (Recall that the connections array is assumed to be sorted by the strength of the connections, with the stronger connections appearing first).

```
SELECT [ connections[$element > 100 AND $pos < 10] ]
AS interestingConnections
FROM users
```

Example 6-56 Array-Filter Step Expression

Count the total number of phones numbers with areacode 650.

```
SELECT count(u.address.phones[$element.area = 650])
FROM users u
```

Example 6-57 Array-Filter Step Expression

To count the total number of people with at least one phone in the 650 areacode, a case expression (see [Case Expressions](#)) and the exists operator (see [Exists Operator](#)) must be used.

```
SELECT count(CASE
WHEN EXISTS u.address.phones[$element.area = 650] THEN 1
ELSE 0
END)
FROM users u
```

Array-Slice Step Expressions

Syntax

`array_slice_step ::= "[" [expression] ":" [expression] "]"`

Semantics

Array slice steps are meant to be used primarily with arrays. In general, an array slice step selects elements of arrays based only on the element positions. The elements to select are the ones whose positions are within a range between a "low" position and a "high" position. The low and high positions are computed by two boundary expressions: a "low" expression for the low position and a "high" expression for the high position. Each boundary expression must return at most one item of type LONG or INTEGER, or NULL. The low and/or the high expression may be missing. The context-item variable (\$) is available during the computation of the boundary expressions.

An array filter step processes each context item as follows:

- If the context item is not an array, an array is created and the context item is added to that array. Then the array filter is applied to this single-item array as described below.
- If the context item is an array, the boundary expressions are computed, if present. If any boundary expression returns NULL or an empty result, the context item is skipped. Otherwise, let L and H be the values returned by the low and high expressions, respectively. If the low expression is absent, L is set to 0. If the high expression is absent, H is set to the size of the array - 1. If L is < 0, L is set to 0. If H > array_size - 1, H is set to array_size - 1. After L and H are computed, the step selects all the elements between positions L and H (L and H included). If L > H no elements are selected.

Notice that based on the above rules, slice steps are actually a special case of filter steps. For example, a slice step with both boundary expressions present, is equivalent to `<input expr>[<low expr> <= $pos and $pos <= <high expr>]`. Slice steps are provided for convenience (and better performance).

Example 6-58 Array-Slice Step Expression

Select the strongest connection of the user with id 10.

```
SELECT connections[0] AS strongestConnection
FROM users
WHERE id = 10
```

Example 6-59 Array-Slice Step Expression

For user 10, select his/her 5 strongest connections (i.e. the first 5 ids in the "connections" array).

```
SELECT [ connections[0:4] ] AS strongConnections
FROM users
WHERE id = 10
```

Notice that the slice expression will return at most 5 ids; if user 10 has fewer than 5 connections, all of his/her connections will be returned.

Example 6-60 Array-Slice Step Expression

For user 10, select his/her 5 weakest connections (i.e. the last 5 ids in the "connections" array).

```
SELECT [ connections[size($) - 5 : ] ] AS weakConnections
FROM users
WHERE id = 10
```

In this example, `size()` is a function that returns the size of a given array, and `$` is the context array, i.e., the array from which the 5 weakest connections are to be selected.

Comparison Expressions

This section describes various comparison expressions in Oracle NoSQL Database.

If you want to follow along with the examples, create the tables as described in the [Tables used in the Examples](#) section.

Logical Operators: AND, OR, and NOT

Syntax

`expression ::= or_expression`

`or_expression ::= and_expression | (or_expression OR and_expression)`

`and_expression ::= not_expression | (and_expression AND not_expression)`

`not_expression ::= [NOT] is\_null\_expression`

Semantics

The binary AND and OR operators and the unary NOT operator have the usual semantics. Their operands are conditional expressions, which must have type `BOOLEAN`. An empty result from an operand is treated as the false value. If an operand returns `NULL`, then:

- The AND operator returns false if the other operand returns false; otherwise, it returns `NULL`.
- The OR operator returns true if the other operand returns true; otherwise it returns `NULL`.
- The NOT operator returns `NULL`.

Example 6-61 Fetch the id, first name, last name, and age for users who are not in the age group of 20 to 40 or whose income is greater than 90K

```
SELECT id, firstName, lastName FROM users WHERE NOT age BETWEEN 20 AND 40 OR
income > 90000
```

Consider an application that maintains the user data. The age field holds the age of the user and the income field includes the income of the user. In the above query, you use a combination of operators to get the list of users who are not in the age group of 20 to 40 years or have an income greater than 90K. A [BETWEEN](#) operator verifies if the input expression is within the range of the boundary values. Since you want users who are either less than 20 years or more than 40 years of age, use the `BETWEEN` operator on the age field with the

logical operator NOT to fetch the users outside the given range. To fetch the list of users with income greater than 90K, use the value comparison operator '>' to compare the income field of the users. You use the OR operator to fetch the list of users who satisfy either of the conditions.

Output:

```
{"id":30,"firstName":"Adam","lastName":"Smith","age":45}
```

Example 6-62 Fetch the list of male passengers from the airline baggage tracking application who are bound toward FRA station and carrying only one checked bag

```
SELECT
fullname,
s.bagInfo[].flightLegs[].fltRouteDest,
ticketNo
FROM BaggageInfo s
WHERE gender = 'M'
AND s.bagInfo[].flightLegs[].fltRouteDest=any "FRA"
AND (size(s.bagInfo) = 1)
```

In the airline baggage tracking application, you can fetch the details of the male passengers who are bound toward a specific destination. The `fltRouteDest` field in the `BaggageInfo` table includes the destination airport code for each travel leg. The `fltRouteDest` field can hold the final station or a transit station. To retrieve a list of male passengers who are bound towards FRA and carrying only one checked bag, you perform value comparison operations and use the AND operator to narrow down the results to successfully match all the criteria. You fetch the list of all the male passengers by comparing the gender field. You retrieve the list of passengers bound towards the FRA station by performing a string comparison operation on the `fltRouteDest` field with the given airport code 'FRA'. You compare the [size](#) of the `bagInfo` array to select the passengers having only one checked bag. You use the AND operator to fetch the list of passengers that satisfy all the conditions mentioned above.

Here, the `s.bagInfo.fltRouteDest` is a sequence. The [value comparison operators](#) cannot operate on sequences of more than one item. A [sequence comparison operator](#) `any` is used in addition to the value comparison operator `=` to compare the `fltRouteDest` fields.

Note

If there is more than one logical operator in the query, ensure the proper usage of parenthesis. The parenthesis is assessed first in the order of evaluation.

Output:

```
{"fullname":"Henry Jenkins","fltRouteDest":
["ORD","FRA"],"ticketNo":176234463813}
{"fullname":"Raymond Griffin","fltRouteDest":
["FRA","HKG"],"ticketNo":1762399766476}
```


IS NULL and IS NOT NULL Operators

Syntax

```
is_null_expression ::= condition_expression [IS [NOT] NULL]
```

```
condition_expression ::=  
    comparison\_expression | exists\_expression  
    | is\_of\_type\_expression | in\_expression
```

Semantics

The IS NULL operator tests whether the result of its input expression is NULL. If the input expression returns more than one item, an error is raised. If the result of the input expression is empty, IS NULL returns false. Otherwise, IS NULL returns true if and only if the single item computed by the input expression is NULL. The IS NOT NULL operator is equivalent to NOT (IS NULL cond_expr). NULL is explained in [Table 2-2](#).

Example 6-63 Select the id, first name, and last name of all users who do not have a known income

Consider an application that maintains the user data. See users table in the [Tables used in the Examples](#) topic.

```
SELECT id, firstName, lastName FROM users u  
WHERE u.income IS NULL
```

Explanation: Assuming that a NULL value is populated in the user table if a user does not have any known income, in the above query, you fetch the details of users whose income field has a NULL value.

Output:

```
{"id":40,"firstName":"Joanna","lastName":"Smith"}
```

Example 6-64 Fetch the ticket number, the full name of passengers from the airline baggage tracking application whose checked bag has a tag associated with it

Consider the airline baggage tracking application. See BaggageInfo table in the [Tables used in the Examples](#) topic.

```
SELECT ticketNo,fullname  
FROM BaggageInfo bag  
WHERE EXISTS bag.bagInfo.tagNum [$element IS NOT NULL]
```

Explanation: In the airline baggage tracking application, there is a unique tag number associated with every checked bag carried by the passenger. In this query, you fetch the details of passengers who have a tag number, which means the tagNum field in the bagInfo table is not null. You use the [EXISTS](#) operator to verify whether or not the tagNum field includes a NULL value.

Output:

```
{"ticketNo":1762330498104,"fullname":"Michelle Payne"}
{"ticketNo":1762340683564,"fullname":"Mary Watson"}
{"ticketNo":1762377974281,"fullname":"Kendal Biddle"}
{"ticketNo":1762320569757,"fullname":"Lucinda Beckman"}
{"ticketNo":1762392135540,"fullname":"Adelaide Willard"}
{"ticketNo":1762399766476,"fullname":"Raymond Griffin"}
{"ticketNo":1762324912391,"fullname":"Elane Lemons"}
{"ticketNo":1762390789239,"fullname":"Zina Christenson"}
{"ticketNo":1762340579411,"fullname":"Zulema Martindale"}
{"ticketNo":1762376407826,"fullname":"Dierdre Amador"}
{"ticketNo":176234463813,"fullname":"Henry Jenkins"}
{"ticketNo":1762311547917,"fullname":"Rosalia Triplett"}
{"ticketNo":1762320369957,"fullname":"Lorenzo Phil"}
{"ticketNo":1762341772625,"fullname":"Gerard Greene"}
{"ticketNo":1762344493810,"fullname":"Adam Phillips"}
{"ticketNo":1762355527825,"fullname":"Doris Martin"}
{"ticketNo":1762383911861,"fullname":"Joanne Diaz"}
{"ticketNo":1762348904343,"fullname":"Omar Harvey"}
{"ticketNo":1762350390409,"fullname":"Fallon Clements"}
{"ticketNo":1762355854464,"fullname":"Lisbeth Wampler"}
{"ticketNo":1762357254392,"fullname":"Teena Colley"}
```

21 rows returned

Value Comparison Operators

Syntax

comparison_expression ::= [concatenate_expression](#)
[(value_comparison_operator | [any_comparison_operator](#)) add_expression]

value_comparison_operator ::= "=" | "!=" | ">" | ">=" | "<" | "<="

Semantics

Value comparison operators are primarily used to compare 2 values, one produced by the left operand and another from the right operand (this is in contrast to the sequence comparisons, defined in the following section which compare two sequences of values). If any operand returns more than one item, an error is raised. If both operands return the empty sequence, the operands are considered equal (so true will be returned if the operator is =, <=, or >=). If only one of the operands returns empty, the result of the comparison is false unless the operator is ! =.

For the remainder of this section, we assume that each operand returns exactly one item. If an operand returns NULL, the result of the comparison expression is also NULL. Otherwise, the result is a boolean value that is computed as follows.

Among atomic items, if the types of the items are not comparable, false is returned. The following rules defined what atomic types are comparable and how the comparison is done in each case.

- A numeric item is comparable with any other numeric item. If an integer or long value is compared to a float or double value, the integer/long will first be cast to float/double. If one

of the operands is a number value, the other operand will first be cast to number (if not a number already).

- A string item is comparable to another string item (using the java `String.compareTo()` method). A string item is also comparable to an enum item. In this case, before the comparison, the string is cast to an enum item in the type of the other enum item. Such a cast is possible only if the enum type contains a token whose string value is equal to the source string. If the cast is successful, the two enum items are then compared as explained below; otherwise, the two items are incomparable and false is returned.
- Two enum items are comparable only if they belong to the same type. If so, the comparison is done on the ordinal numbers of the two enums (not their string values). As mentioned above, an enum item is also comparable to a string item, by casting the string to an enum item.
- Binary and fixed binary items are comparable with each other for equality only. The 2 values are equal if their byte sequences have the same length and are equal byte-per-byte.
- A boolean item is comparable with another boolean item, using the java `Boolean.compareTo()` method.
- A timestamp item is comparable to another timestamp item, even if their precisions are different.
- JNULL (json null) is comparable with JNULL. If the comparison operator is `!=`, JNULL is also comparable with every other kind of item, and the result of such a comparison is always true, except when the other item is also JNULL.

The semantics of comparisons among complex items are defined in a recursive fashion. Specifically:

- A record is comparable with another record for equality only and only if they contain comparable values. More specifically, to be equal, the 2 records must have equal sizes (number of fields) and for each field in the first record, there must exist a field in the other record such that the two fields are at the same position within their containing records, have equal field names, and equal values.
- A map is comparable with another map for equality only and only if they contain comparable values. Remember that json documents are modelled as maps, so 2 json documents can be compared for equality. More specifically, to be equal, the 2 maps must have equal sizes (number of fields) and for each field in the first map, there must exist a field in the other map such that the two fields have equal names and equal values.
- An array is comparable to another array if the elements of the 2 arrays are comparable pair-wise. Comparison between 2 arrays is done lexicographically, that is, the arrays are compared like strings, with the array elements playing the role of the "characters" to compare.

As with atomic items, if two complex items are not comparable according to the above rules, false is returned. Furthermore, comparisons between atomic and complex items return false always.

The reason for returning false for incomparable items, instead of raising an error, is to handle truly schemaless applications, where different table rows may contain very different data or differently shaped data. As a result, even the writer of the query may not know what kind of items an operand may return and an operand may indeed return different kinds of items from different rows. Nevertheless, when the query writer compares "something" with, say, an integer, they expect that the "something" will be an integer and they would like to see results from the table rows that fulfill that expectation, instead of the whole query being rejected because some rows do not fulfill the expectation.

Example 6-65 Fetch the passenger details from the airline baggage tracking application who have more than two transits and did not board from the SFO station

```
SELECT
fullname,
s.bagInfo.routing
FROM BaggageInfo s
WHERE (size(s.bagInfo[0].flightLegs) >= 3)
AND s.bagInfo[0].flightLegs[0].fltRouteSrc !=any "SFO"
```

Explanation: In the airline baggage tracking application, you can fetch the list of passengers who have more than two transits in their journey. The bagInfo array in the BaggageInfo table contains the information on the checked bags for passengers. The flightLegs array in the bagInfo JSON includes the source and transit details with each record corresponding to a travel leg. The first record of the flightLegs array, that is, flightLegs[0] has the details of the source location. In the above query, you fetch the details of the passengers who have more than two transits, which means there must be at least three records in the flightLegs array including the source location. You compare the size of the flightLegs array using the '>=' operator. You also want to filter out the passengers who did not originate from SFO airport, so you use a != operator here. It is possible that the passengers have more than one checked bags, in which case there will be more than one element in the bagInfo array. You must consider only the first element of the bagInfo array, that is, bagInfo[0] during value comparison to avoid duplication of results.

Since the flightLegs is an array, the left operand of the comparison operator != is a sequence with more than one item. Hence, use the sequence comparison operator any in addition to the value comparison operator.

Output:

```
{"fullname":"Fallon Clements","routing":"MXP/CDG/SLC/BZN"}
{"fullname":"Elane Lemons","routing":"MXP/CDG/SLC/BZN"}
{"fullname":"Doris Martin","routing":"BZN/SEA/CDG/MXP"}
```

JSON collection table:

The following example applies a value comparison operator on a JSON collection table. Consider a sample row from the JSON collection table created for a [shopping application](#):

```
{"contactPhone":"1917113999","address":{"city":"San
Jose","number":501,"state":"San
Francisco","street":"Maine","zip":95095},"cart":
[{"item":"wallet","priceperunit":950,"quantity":2},{ "item":"wall
art","priceperunit":9500,"quantity":1}], "firstName":"Sharon","gender":"F","las
tName":"Willard","notify":"yes","wishlist":
[{"item":"Tshirt","priceperunit":500},{ "item":"Jenga","priceperunit":850}]}
```

Example 6-66 Fetch from the storeAcct table the details of shoppers who have wishlisted the item Jenga.

```
SELECT contactPhone, firstName
FROM storeAcct s
WHERE EXISTS s.wishlist[$element.item ="Jenga"]
```

Explanation: In the storeAcct table, the items wishlisted by the shoppers are stored in the JSON array wishlist. To fetch the details of the shoppers with Jenga as a wishlisted item, you verify whether or not an item by the specified name exists in the wishlist field. Use the operator '=' to match the item.

Output:

```
{
  "contactPhone" : "1917113999",
  "firstName" : "Sharon"
}
```

Sequence Comparison Operators

Syntax

```
any_comparison_operator ::= "=any" | "!=any" | ">any" | ">=any" | "<any" |
"<=any"
```

Semantics

Comparisons between two sequences is done via another set of operators: =any, !=any, >any, >=any, <any, <=any. These any operators have existential semantics: the result of an any operator on two input sequences S1 and S2 is true if and only if there is a pair of items i1 and i2, where i1 belongs to S1, i2 belongs to S2, and i1 and i2 compare true via the corresponding value comparison operator. Otherwise, if any of the input sequences contains NULL, the result is NULL. Otherwise, the result is false.

Example 6-67 Fetch from the airline baggage tracking data, the passenger name where the flight time is greater than 2019-02-01T16:00:00Z

```
SELECT fullname, $flights.flightDate
FROM BaggageInfo bag, bag.bagInfo[0].flightLegs AS $flights
WHERE $flights.flightDate >any "2019-02-01T16:00:00Z"
```

Explanation: You fetch the full name and the flight dates of all passengers whose departure time is greater than the given value. Here the operand on the left hand of the ">" operator, that is, bag.bagInfo[0].flightLegs.flightDate is a sequence of values. If you try using the regular comparison operator instead of the sequence operator, you get an error as shown below:

```
Error handling command SELECT fullname, $flights.flightDate
FROM BaggageInfo bag, bag.bagInfo[0].flightLegs AS $flights
WHERE $flights.flightDate > "2019-02-01T16:00:00Z": Error: at (3, 6) The left
operand of comparison operator > is a sequence with more than one items.
Comparison operators cannot operate on sequences of more than one items.
```

Instead, you must use a sequence operator any along with the value comparison operator ">" here.

You use only the first element of the bagInfo JSON array to avoid duplication of results for passengers carrying more than one bag. The departure time for the transit flights is also included in the output. Here, you can compare the string-formatted dates in ISO-8601 format

due to the natural sorting order of strings without having to cast them into timestamp data types.

Output:

```
{ "fullname": "Michelle Payne", "flightDate":
["2019-02-02T12:00:00Z", "2019-02-02T13:10:00Z", "2019-02-02T12:10:00Z"]}
{ "fullname": "Zulema Martindale", "flightDate":
["2019-02-25T10:00:00Z", "2019-02-25T10:23:00Z"]}
{ "fullname": "Zina Christenson", "flightDate":
["2019-02-04T00:00:00Z", "2019-02-04T00:12:00Z"]}
{ "fullname": "Elane Lemons", "flightDate":
["2019-03-15T08:00:00Z", "2019-03-15T08:13:00Z", "2019-03-15T08:13:00Z"]}
{ "fullname": "Omar Harvey", "flightDate":
["2019-03-01T12:00:00Z", "2019-03-01T12:20:00Z"]}
{ "fullname": "Henry Jenkins", "flightDate":
["2019-03-02T05:00:00Z", "2019-03-02T05:24:00Z"]}
{ "fullname": "Rosalia Triplett", "flightDate":
["2019-02-11T01:00:00Z", "2019-02-11T02:12:00Z"]}
{ "fullname": "Kendal Biddle", "flightDate":
["2019-03-04T06:00:00Z", "2019-03-04T07:02:00Z"]}
{ "fullname": "Mary Watson", "flightDate":
["2019-03-13T00:00:00Z", "2019-03-13T00:21:00Z"]}
{ "fullname": "Gerard Greene", "flightDate":
["2019-03-07T04:00:00Z", "2019-03-07T05:08:00Z", "2019-03-07T04:08:00Z"]}
{ "fullname": "Joanne Diaz", "flightDate":
["2019-02-16T06:00:00Z", "2019-02-16T06:24:00Z"]}
{ "fullname": "Adelaide Willard", "flightDate":
["2019-02-15T01:00:00Z", "2019-02-15T01:21:00Z"]}
{ "fullname": "Fallon Clements", "flightDate":
["2019-02-21T12:00:00Z", "2019-02-21T12:19:00Z", "2019-02-21T12:19:00Z"]}
{ "fullname": "Teena Colley", "flightDate":
["2019-02-13T07:00:00Z", "2019-02-13T07:17:00Z"]}
{ "fullname": "Lorenzo Phil", "flightDate":
["2019-03-12T03:00:00Z", "2019-03-12T10:11:00Z", "2019-03-12T14:11:00Z"]}
{ "fullname": "Lucinda Beckman", "flightDate":
["2019-03-12T03:00:00Z", "2019-03-12T04:11:00Z", "2019-03-12T03:11:00Z"]}
{ "fullname": "Doris Martin", "flightDate":
["2019-03-22T07:00:00Z", "2019-03-22T07:23:00Z", "2019-03-22T07:23:00Z"]}
{ "fullname": "Raymond Griffin", "flightDate":
["2019-02-03T04:00:00Z", "2019-02-03T04:21:00Z"]}
{ "fullname": "Lisbeth Wampler", "flightDate":
["2019-02-09T06:00:00Z", "2019-02-09T20:01:00Z"]}
{ "fullname": "Dierdre Amador", "flightDate": "2019-03-07T07:00:00Z" }
```

Example 6-68 Select the id, lastName and address for users who are connected with the user with id 3.

```
SELECT id, lastName, address FROM users
WHERE connections[] =any 3
```

Notice the use of [] after connections: it is an array filter step (see [Array-Filter Step Expressions](#)), which returns all the elements of the connections array as a sequence (it is unnesting the array).

Example 6-69 Select the id and lastName for users who are connected with any users having id greater than 100

```
SELECT id, lastName FROM users
WHERE connections[] >any 100
```

Example 6-70 Select the id of each user who is connected with a user having id greater than 10 and is also connected with a user having id less than 100

```
SELECT id FROM users u
WHERE 10 <any u.connections[]
      AND u.connections[] <any 100
```

Notice that the above query is not the same as the query: *select the id of each user who is connected with a user having an id in the range between 10 and 100*. In the first query, you are looking for some connection with id greater than 10 and another connection (which may or may not be the same as the 1st one) with id less than 100. In the second query you are looking for some connection whose id is between 10 and 100. To make the difference clear, consider a Users table with only 2 users (say with ids 200 and 500) having the following connections arrays respectively: [1, 3, 110, 120] and [1, 50, 130]. Both of these arrays satisfy the predicates in the first query, and as a result, both users will be selected. On the other hand, the second query will not select user 200, because the array [1, 3, 110, 120] does not contain any element in the range 10 to 100.

The second query can be written by a combination of an EXISTS operator and an array filtering step:

```
SELECT id FROM users u
WHERE EXISTS u.connections
      [10 < $element AND $element < 100]
```

and the first query, with the 2 <any operators, is equivalent to the following one:

```
SELECT id FROM users u
WHERE EXISTS u.connections[10 < $element]
      AND EXISTS u.connections[$element < 100]
```

Example 6-71 Select the first and last name of all users who have a phone number with area code 650

Notice that although we could have used [] after phones in this query, it is not necessary to do so, because the phones array (if it is indeed an array) is unnested implicitly by the .area step that follows.

```
SELECT firstName, lastName FROM users u
WHERE u.address.phones.area =any 650
```

JSON collection table:

The following example applies the sequence comparison operator on a JSON collection table. Consider a sample row from the JSON collection table created for a [shopping application](#):

```
{"contactPhone":"1517113582","address":
{"city":"Houston","number":651,"state":"TX","street":"Tex
```

```
Ave", "zip": 95085}, {"cart": null, "firstName": "Dierdre", "lastName": "Amador", "orders": [{"EstDelivery": "2023-11-01", "item": "handbag", "orderId": "201200", "priceperunit": 350}, {"EstDelivery": "2023-11-01", "item": "Lego", "orderId": "201201", "priceperunit": 5500}]}
```

Example 6-72 Fetch the details from shoppers who have purchased a handbag and the stipulated delivery is after October 31st, 2023

```
SELECT contactPhone, firstName
FROM storeAcct s
WHERE s.orders[].item =any "handbag" AND s.orders[].EstDelivery>=any
"2023-10-31"
```

Explanation: To fetch the details from shoppers who have purchased a handbag that is expected to be delivered after October 31st, you compare the `item` and `EstDelivery` fields with the required values using the sequence comparison operator `any`. You use the [AND](#) operator to fetch the rows that match both conditions.

Here, you can compare the `EstDelivery` without casting into a timestamp data type as it is a string-formatted date in ISO-8601 format and the natural sorting order of strings applies.

Output:

```
{
  "contactPhone" : "1517113582",
  "firstName" : "Dierdre"
}
```

BETWEEN Operator

Syntax

```
between_expression ::= input_expression BETWEEN low_bound_expression AND
                        high_bound_expression
input_expression   := concatenate\_expression
low_bound_expression := concatenate\_expression
high_bound_expression := concatenate\_expression
```

Semantics

The BETWEEN operator checks if the input expression value is in between the lower and the higher expressions (including the boundary values). This is equivalent to:

```
low_bound_expression <= input_expression AND input_expression <=
high_bound_expression
```

The BETWEEN operator internally performs the following:

1. Two value comparison operations: It checks the following conditions,
 - a. If the low bound expression is `<=` the input expression.

- b. If the input expression is $\leq$ the high bound expression.
- 2. Logical operation - the logical operator AND is applied to the results.

This essentially verifies if the input expression is within the range of the boundary values. The operation returns a TRUE value if both the expressions return TRUE. The operation returns a FALSE value if any one of the expressions returns FALSE. If either of the expressions is NULL or leads to a NULL value, the result of the operation is also NULL. If any expression returns more than one item, an error is raised as the comparison operators do not operate on sequences of more than one item. For more details on the value comparison operators and logical operators, see [Comparison Expressions](#).

See the semantics in the [Value Comparison Operators](#) topic to understand the comparison of the input expressions for different data types.

Example 6-73 Fetch the passenger details whose ticket numbers are in a certain range.

```
SELECT
fullname,
ticketNo
FROM baggageinfo s
WHERE ticketNo BETWEEN 17623000000000 and 17623500000000
```

Explanation: In this query, you fetch the details of passengers whose ticket numbers are between 17623000000000 and 17623500000000 from the baggageInfo table. You use the BETWEEN operator in the WHERE clause to select and display the rows that fall within the required range.

Output:

```
{"fullname":"Lorenzo Phil","ticketNo":1762320369957}
{"fullname":"Elane Lemons","ticketNo":1762324912391}
{"fullname":"Michelle Payne","ticketNo":1762330498104}
{"fullname":"Lucinda Beckman","ticketNo":1762320569757}
{"fullname":"Rosalia Triplett","ticketNo":1762311547917}
{"fullname":"Zulema Martindale","ticketNo":1762340579411}
{"fullname":"Omar Harvey","ticketNo":1762348904343}
{"fullname":"Adam Phillips","ticketNo":1762344493810}
{"fullname":"Mary Watson","ticketNo":1762340683564}
{"fullname":"Gerard Greene","ticketNo":1762341772625}
```

Example 6-74 Fetch the passenger details and routing information of the baggage that falls within a range of reservation codes.

```
SELECT fullname AS FULLNAME,
confNo AS RESERVATION,
s.bagInfo.routing AS ROUTINGINFO
FROM BaggageInfo s
WHERE confNo BETWEEN 'LE6J4Z' and 'ZG8Z5N'
ORDER BY confNo
```

Explanation: Every passenger has a reservation code (confNo). In this query, you fetch the passenger details, reservation code, and routing details for the baggage whose reservation codes are within the range of LE6J4Z and ZG8Z5N. You use the BETWEEN operator in the WHERE clause to perform a string comparison of the confNo value with the lower and the

upper boundary values in the input strings. Only the rows that are within the range are selected and displayed in the output.

Output:

```
{ "FULLNAME": "Adam Phillips", "RESERVATION": "LE6J4Z", "ROUTINGINFO": "MIA/LAX/MEL" }
{ "FULLNAME": "Elane Lemons", "RESERVATION": "LN0C8R", "ROUTINGINFO": "MXP/CDG/SLC/BZN" }
{ "FULLNAME": "Gerard Greene", "RESERVATION": "MC0E7R", "ROUTINGINFO": "SFO/IST/ATH/JTR" }
{ "FULLNAME": "Henry Jenkins", "RESERVATION": "MZ2S5R", "ROUTINGINFO": "SFO/ORD/FRA" }
{ "FULLNAME": "Omar Harvey", "RESERVATION": "OH2F8U", "ROUTINGINFO": "MEL/LAX/MIA" }
{ "FULLNAME": "Kendal Biddle", "RESERVATION": "PQ1M8N", "ROUTINGINFO": "JFK/IST/VIE" }
{ "FULLNAME": "Zina Christenson", "RESERVATION": "QB100J", "ROUTINGINFO": "MIA/LAX/MEL" }
{ "FULLNAME": "Lorenzo Phil", "RESERVATION": "QI3V6Q", "ROUTINGINFO": ["SFO/IST/ATH/JTR", "SFO/IST/ATH/JTR"] }
{ "FULLNAME": "Lucinda Beckman", "RESERVATION": "QI3V6Q", "ROUTINGINFO": "SFO/IST/ATH/JTR" }
{ "FULLNAME": "Michelle Payne", "RESERVATION": "RL3J4Q", "ROUTINGINFO": "SFO/IST/ATH/JTR" }
{ "FULLNAME": "Teena Colley", "RESERVATION": "TX1P7E", "ROUTINGINFO": "MSQ/FRA/HKG" }
{ "FULLNAME": "Fallon Clements", "RESERVATION": "XT107T", "ROUTINGINFO": "MXP/CDG/SLC/BZN" }
{ "FULLNAME": "Raymond Griffin", "RESERVATION": "XT6K7M", "ROUTINGINFO": "MSQ/FRA/HKG" }
{ "FULLNAME": "Dierdre Amador", "RESERVATION": "ZG8Z5N", "ROUTINGINFO": "JFK/MAD" }
```

Example 6-75 Find the passengers who traveled from MIA within a fortnight from 15th Feb 2019.

```
SELECT fullname,
FROM BaggageInfo bag
WHERE exists bag.bagInfo.flightLegs[$element.fltRouteSrc = "MIA"
AND
$element.flightDate BETWEEN "2019-02-15T00:00:00Z" and "2019-03-02T00:00:00Z"]
```

Explanation: In this query, you fetch the details of the passengers who traveled from MIA between the 15th of Feb 2019 and the 2nd of March 2019. The flightDate field within the bagInfo JSON field contains the travel dates to the destination points. You use the BETWEEN operator to compare the flightDate in the passenger data with the upper and the lower range of the specified dates. The flightDate is a string and is directly compared with the supplied dates, which are also string values. You narrow down the passenger records listed within this range further to include only MIA as the source station using the AND operator. Here the flight source could be the starting point of the flight or any transit airport.

Output:

```
{ "fullname": "Zulema Martindale" }
{ "fullname": "Joanne Diaz" }
```

IN Operator

Syntax

```
in_expression ::= in1_expression | in2_expression |
                 in3_expression | in4_expression

in1_expression ::= "(" concatenate_expression
                  ("," concatenate_expression)* ")"
                 IN "(" expression ("," expression)* ")"
in2_expression ::= concatenate_expression
                 IN "(" expression ("," expression)* ")"
in3_expression ::= concatenate_expression IN path_expression
in4_expression ::= "(" concatenate_expression
                  ("," concatenate_expression)* ")" IN path_expression
```

Semantics

The IN operator is essentially a compact alternative to a number of OR-ed equality conditions.

Example 6-76 Fetch users who are 22, 25, or 43 years old

```
SELECT * FROM users WHERE age IN (22, 25, 43)
```

Output:

```
{"id":10,"firstName":"John","lastName":"Smith","otherNames":
[{"first":"Johnny","last":"Good"}, {"first":"Johnny2","last":"Brave"},
{"first":"Johnny3","last":"Kind"},
{"first":"Johnny4","last":"Humble"}], "age":22, "income":45000, "address":
{"city":"Santa Cruz", "number":101, "phones":
[{"area":408, "kind":"work", "number":4538955},
{"area":831, "kind":"home", "number":7533341},
{"area":831, "kind":"mobile", "number":7533382}], "state":"CA", "street":"Pacific
Ave", "zip":95008}, "connections":[30,55,43], "hobbies":
["Reading", "Soccer", "Hiking", "Gardening"]}
```

The query above is equivalent to

```
SELECT * FROM users WHERE age = 22 OR age = 25 OR age = 43
```

and the query,

```
SELECT * FROM users
  WHERE (firstName, lastName) IN
        (("John", "Smith"), ("Peter", "Paul"), ("Mary", "Ann"))
```

is equivalent to

```
SELECT * FROM users
  WHERE (firstName = "John" AND lastName = "Smith") OR
```

```
(firstName = "Peter" AND lastName = "Paul") OR
(firstName = "Mary" AND lastName = "Ann")
```

As shown in the grammar, there are 4 syntactic variants of the IN operator. The `in1_expression` and `in2_expression` follow the standard SQL syntax. The `in2_expression` one is actually a special case of the `in1_expression`, for the case when there is only one expression in the left-hand-side of the operator. For the `in1_expression`, if *K* is the number of expressions in the left-hand-side, then each expression list in the right-hand-side must consist of *K* expressions. If *N* is the number of expression lists in the right-hand-side, then the whole IN condition is equivalent to:

```
(expr1 = expr11 and expr2 = expr12 and exprK = expr1K) or
(expr1 = expr21 and expr2 = expr22 and exprK = expr2K) or
... or
(expr1 = exprN1 and expr2 = exprN2 and exprK = exprNK)
```

However, in addition to being more compact, queries using IN operators will be executed more efficiently if appropriate indexes exist. For example, if table `users` has an index on columns `age`, `firstName` and `lastName`, then both of the above IN queries will use that index to find the qualifying rows, whereas the equivalent OR queries will be executed via full table scans. See also examples in [Finding Applicable Indexes](#).

The `in3_expression` and `in4_expression` variants allow a relative large number of search keys to be provided via a single bind variable. For example, if the `$keys` variable is bound to the array `["John", "Smith", "Peter", "Paul", "Mary", "Ann"]`, then the following query is equivalent to the second IN query above.

```
DECLARE $keys ARRAY(json);
SELECT * FROM users
       WHERE (firstName, lastName) IN $keys[]
```

In general, with the `in3_expression` and `in4_expression` variants, the expression in the right-hand-side is evaluated first. If the number *M* of items in the resulting sequence is less than the number *K* of expressions in the left-hand-side, the result of the IN operator is false. If *M* is not a multiple of *K*, the last (*M mod K*) items in the sequence are discarded and *M* is set to the number of remaining items. Then, the IN expression is equivalent to:

```
(expr1 = k1 and expr2 = k2 and exprK = kK) or
(expr1 = kK+1 and expr2 = kK+2 and exprK = k2*K) or
... or
(expr1 = kM-K and expr2 = kM-K+1 and exprK = kM)
```

However, an additional type-checking restriction applies in this case: in each of the above equality conditions, the type of the right-hand-side item must be a subtype of the left-hand-side type.

Regular Expression Conditions

The `regex_like` function performs regular expression matching. A regular expression is a pattern that the regular expression engine attempts to match with an input string. The syntax for invoking the `regex_like` function in a query is the same as all other functions, described in the [Function Calls](#) section. The `regex_like` function has two signatures with two and three parameters, respectively.

Syntax

```
boolean regex_like(any*, string)
```

```
boolean regex_like(any*, string, string)
```

Semantics

The `regex_like` function provides functionality similar to the `LIKE` operator in standard SQL, that is, it can be used to check if an input string matches a given pattern. The input string and the pattern are computed by the first and second arguments, respectively. A third, optional, argument specifies a set of flags that affect how the matching is done.

Normally, the `regex_like` function expects each of its arguments to return a single string. If that is not the case, it behaves as follows:

- If it can be detected at compile time that the first argument will never return a string, it raises a compile-time error. Otherwise, it returns false if the first argument returns nothing, more than one item, or a single item that is neither a string nor NULL.
- It raises an error if the pattern or flags do not return a single string or NULL.
- It returns NULL if any of the arguments returns a single NULL.

Otherwise, the `regex_like` function behaves as follows:

- Raises an error if the pattern string is not valid or its length is greater than 512 characters.
- Returns false if the pattern does not match the input string.
- Returns true if the pattern matches the input string.

The pattern string is the regular expression against which the input text is matched. The syntax of the pattern string is a subset of the one supported by the Java Pattern class. Specifically, each character in a regular expression is either a literal character that matches itself (for example, the pattern string `x` matches the character `'x'`), or a meta character, that specifies a *construct* having a special meaning. Only the following metacharacters are supported:

- Period (`.`)
You use period metacharacter to match every character except a new line
- Greedy quantifier (`*`)
You use the greedy quantifier (`*`) to indicate zero or more occurrences of the preceding element.

For example, the character `D` with a combination of period (`.`) and (`*`) `"D.*"`, matches any string that starts with the character `'D'` and is followed by zero or more characters.

- Quotation constructs (`\`, `\Q`, `\E`)
You use the backslash `'\'` character as a quotation construct with other metacharacters to instruct the regular expression engine to interpret the following metacharacter as a literal character.

For example, you use the `"\*"` pattern to match the literal asterisk `'*'` character and not interpret it as the greedy quantifier (`*`).

You use `\Q` and `\E` quotation constructs in the query to match the exact string pattern that starts after the `\Q` character and ends at the `\E` character.

For example, you use `"\Q**Houston**\E"` to match any string that has a value `'**Houston**'`.

Note

In the SQL shell or other equivalent tools, you must use a double backslash '\\' in place of a single backslash '\' quotation construct.

- **Escape sequences**
You use the escape sequences to match certain predefined characters. The following escape sequences are supported:

The flags string contains one or more characters, where each character is a flag specifying some particular behavior. The full list of acceptable characters and their semantics is listed in the following table:

Table 6-2 Escape sequences

Character	Description
\x{h...h}	Matches the character with hexadecimal value 0xh...h (Character.MIN_CODE_POINT <= 0xh...h <= Character.MAX_CODE_POINT)
\xhh	Matches the character with hexadecimal value 0xhh
\uhhhh	Matches the Unicode character specified by the hexadecimal number 0xhhhh. Example: 'u000A' matches the newline character
\t	Matches the tab character ('\u0009')
\r	Matches the carriage-return character ('\u000D')
\n	Matches the newline (line feed) character ('\u000A')
\f	Matches the form-feed character ('\u000C')
\e	Matches the escape character ('\u001B')
\cx	Matches the control character corresponding to x
\a	Matches the alert (bell) character ('\u0007')
\\	Matches the backslash '\' character ('\u005C')
\0nn	Matches the character with octal value 0nn (0 <= n <= 7)
\0n	Matches the character with octal value 0n (0 <= n <= 7)
\0mnn	Matches the character with octal value 0mnn (0 <= m <= 3, 0 <= n <= 7)

The flag contains one or more characters, where each character flag specifies some particular behavior. The full list of flag characters and their semantics are listed in the following table:

Table 6-3 Predefined Quoted Characters

Flag	Full Name	Description
"d"	UNIX_LINES	Enables Unix lines mode. The new line character ('\n') is the only line termination method recognized in this mode.
"i"	CASE_INSENSITIVE	Enables case-insensitive matching. By default, CASE_INSENSITIVE matching assumes that only characters in the US-ASCII character set are being matched. You can enable Unicode-aware CASE_INSENSITIVE by specifying the UNICODE_CASE flag in conjunction with this flag. Specifying this flag may impose a slight performance penalty.

Table 6-3 (Cont.) Predefined Quoted Characters

Flag	Full Name	Description
"x"	COMMENTS	Permits white space and comments in pattern. In this mode, white space is ignored, and embedded comments starting with # are ignored until the end of a line.
"l"	LITERAL	When LITERAL is specified then the input string that specifies the pattern is treated as a sequence of literal characters. There is no special meaning for Metacharacters or escape sequences. The flags CASE_INSENSITIVE and UNICODE_CASE retain their impact on matching when used in conjunction with this flag. The other flags become superfluous.
"s"	DOTALL	Enables DOTALL mode. In DOTALL mode, the expression dot (.) matches any character, including a line terminator. However, by default, the expression dot (.) does not match line terminators.
"u"	UNICODE_CASE	When you enable the CASE_INSENSITIVE flag, by default, it matches using only the characters in the US-ASCII character set. When you specify UNICODE_CASE then it performs CASE_INSENSITIVE matching using the Unicode standard. Specifying this flag may impose a performance penalty.
"c"	CANON_EQ	When this flag is specified then two characters will be considered to match if, and only if, their full canonical decompositions match. When you specify this flag, the expression "a\u030A", for example, will match the string "\u00E5. By default, matching does not take canonical equivalence into account. Specifying this flag may impose a performance penalty.
"U"	UNICODE_CHARACTER_CLASS	Enables the Unicode version of Predefined character classes and POSIX character classes. When you specify this flag, then the (US-ASCII only) Predefined character classes and POSIX character classes are in conformance with Unicode Technical Standards. See http://unicode.org/reports/tr18/#Compatibility_Properties . The flag implies UNICODE_CASE; it enables Unicode-aware case folding. Specifying this flag may impose a performance penalty.

Note

The `regex_like` function will not be used for index scan optimization.

If you want to follow along with the examples, create and load data into [user data application](#) table illustrated in this section and *airline baggage tracking application* table illustrated in the [Tables used in the Examples](#) topic.

Note

Run the queries provided in the below examples from the SQL shell. Make sure that you use a double backslash `\\` in place of a single backslash `\` for quotation constructs as shown in the queries.

Example 6-77 Fetch from the user data application table the list of users whose last name starts with 'S'

```
SELECT id, firstName, lastName FROM users WHERE regex_like(lastName, "S.*")
```

Explanation: In the above query, you use the `regex_like` function to match the `lastName` field beginning with an 'S' character. You create a pattern with the first character as 'S' followed by the period metacharacter (`.`) and the greedy quantifier (`*`) to match zero or more occurrences of any other character.

Output:

```
{ "id": 10, "firstName": "John", "lastName": "Smith" }
{ "id": 20, "firstName": "Jane", "lastName": "Smith" }
{ "id": 30, "firstName": "Adam", "lastName": "Smith" }
{ "id": 40, "firstName": "Joanna", "lastName": "Smith" }
```

Example 6-78 Fetch from the user data application table the list of users whose last name has at least one 'w' or 'W'

```
SELECT id, lastName FROM users WHERE regex_like(lastname, ".*w.*", "i")
```

Explanation: In the above query, use the `regex_like` function with the required pattern and the 'i' flag to enable matching that is not case-sensitive.

Output:

```
{"id":60,"lastName":"Law"}
{"id":50,"lastName":"Winslet"}
```

Example 6-79 Validate the format of an e-mail address in the user data application table

```
SELECT id, firstName, lastName, email FROM users WHERE
    regex_like(email, ".*@.*\\.\\..*")
```

Explanation: In the above query, you fetch the list of users with a valid email format. In the user data application table, the email field contains the e-mail address of the user. Assuming the email field has an '@' character and ends with a '.string' pattern such as '.com', '.us', '.in', and so forth, you can validate the e-mail address using the `regex_like` function as follows:

- Use the '@' character preceded and followed by a combination of period metacharacter (.) and the greedy quantifier (*). This combination creates a pattern to match zero or more occurrences of other characters before and after the '@' character in the email field.
- To achieve a domain name pattern such as '.com', you use the quotation construct backslash '\' character to match the period as a literal '.' character and not a period (.) metacharacter. Further, use the combination of period metacharacter (.) and the greedy quantifier (*) to allow any domain name.

Note

This example only validates the e-mail address format by checking the pattern availability of '@' and '.' characters in the specified order. It does not assure the validity of the e-mail address itself. All the rows that include the mentioned character pattern are fetched.

Output:

```
{"id":50,"firstName":"Aubrey","lastName":"Winslet","email":"reachaubrey@somemail.co.us"}
{"id":60,"firstName":"Jimmy","lastName":"Law","email":"reachjimlaw@gotmail.co.us"}
{"id":20,"firstName":"Jane","lastName":"Smith","email":"jane.smith201@reachmail.com"}
{"id":10,"firstName":"John","lastName":"Smith","email":"john.smith@reachmail.com"}
```

Example 6-80 Fetch from the user data application table the list of users with a five-star rating for community service

```
SELECT id, firstName, lastName FROM users WHERE
       regex_like(communityService,"\\Q*****\\E")
```

Explanation: In the above query, you fetch the list of users actively involved in community service. Assuming each user is rated for their involvement in community service with a certain '*' rating, you retrieve the list of users from the user data application table who have five stars, that is, "*****" in the communityService field. You use the regex_like function with the quotation constructs '\Q' and '\E' to match the pattern of five asterisk '*' characters. In this query, the asterisk (*) character is used as a literal '*' character and not as a greedy quantifier (*).

Note

If you create the pattern without quotation constructs, an error is generated indicating that the specified pattern for the regex_like function is invalid.

Output:

```
{"id":20,"firstName":"Jane","lastName":"Smith"}
```

Example 6-81 Fetch from the user data application table the list of users whose street attribute of the address field has a suspected data entry error due to the presence of a tab character

```
SELECT id, firstName, users.address.street FROM users WHERE
       regex_like(users.address.street, ".*\t.*")
```

Explanation: You can use the regex_like function to identify the rows that include escape sequences. For a detailed list of supported predefined characters, see [Table 6-2](#). In this query,

you fetch the list of users from the user data application table whose street attribute of the address field erroneously includes a tab character. You use the `regex_like` function with the following pattern:

- Use escape sequence `'\t'` to identify the tab character.
- Use the combination of period metacharacter `(.)` and the greedy quantifier `(*)` before and after the escape sequence to allow zero or more occurrences of any other character in the street attribute.

Output:

```
{"id":70,"firstName":"Dierdre","street":"Maine\t(Suburb)"}
```

Alternatively, you can match the escape sequences using the corresponding Unicode hexadecimal character as shown in the query below. The Unicode value `0x0009` corresponds to a horizontal tab or a tab character. You get the same output as above.

```
SELECT id, firstName, users.address.street FROM users WHERE
regex_like(users.address.street,
           ".*\u0009.*")
```

Note

Oracle NoSQL Database supports insertion of control characters (ASCII code 0~31) and characters with ASCII code > 128 using their Unicode hexadecimal values in the SQL shell or equivalent tools.

For example, you can insert the Escape (ESC) character using its Unicode hexadecimal value `0x001B` as given in the [user data application table](#) (see the state field in the row with `id = 70`).

If you want to fetch the corresponding row, match the pattern through regular expression using `'\e'` as provided in the query below:

```
SELECT id, firstName, users.address[].state FROM users WHERE
       EXISTS(users.address[regex_like($element.state, ".*\\e")])
```

Explanation: In this query, you check whether or not the state attribute of the address field includes a string that ends with an escape character. You use the quotation construct `'\'` to match the escape character preceded by the combination of period metacharacter `(.)` and greedy quantifier `(*)` to allow zero or more occurrences of other characters before the escape character.

Output:

```
{"id":70,"firstName":"Dierdre","state":"TX\u001B"}
```

Example 6-82 Find all the bags that traveled through ORD airport as an intermediate hop between the source of a trip and the final destination

```
SELECT
ticketNo AS TICKETNUM,
```

```
fullName AS NAME,  
BaggageInfo.bagInfo.routing[] AS ROUTING  
FROM BaggageInfo WHERE regex_like(BaggageInfo.bagInfo.routing[], ".*ORD/.*")
```

Explanation: In an airline baggage tracking application, you can fetch the list of passengers who had to transit at certain airports. For each passenger, the routing field in the BaggageInfo table contains the airport codes in the format source/transit/destination. In this query, you use the regex_like function to match the airport code of the transit airport as follows:

- Use the combination of the period (.) metacharacter and greedy quantifier (*) to allow any source airport code characters.
- Include the airport code of the transit airport ORD between the two forward slash '/' characters.
- Use the combination of the period (.) metacharacter and greedy quantifier (*) to allow any destination airport code characters.

Output:

```
{"TICKETNUM":176234463813,"NAME":"Henry Jenkins","ROUTING":"SFO/ORD/FRA"}  
{"TICKETNUM":1762392135540,"NAME":"Adelaide Willard","ROUTING":"GRU/ORD/SEA"}
```

Example 6-83 Find all the passengers with area code 364 in their contact phone.

```
SELECT ticketNo, contactPhone, fullName FROM BaggageInfo WHERE  
regex_like(contactPhone, "364-.*")
```

Explanation: In an airline baggage tracking application, you can fetch the list of passengers having the same area code in their contact phones. The contactPhone field contains the US-based contact number in a three-digit area code followed by a seven-digit local number format. Assuming the contact phone pattern in the contactPhone field is XXX-YYY-ZZZZ where X, Y, and Z are digits between 0-9, you use the regex_like function as follows:

- Use the area code 364 followed by a hyphen '-' character.
- Use the combination of the period (.) metacharacter and greedy quantifier (*) to allow the pattern matching of any three-digit number followed by a hyphen '-' character and a four-digit number.

Output:

```
{"ticketNo":1762320369957,"contactPhone":"364-610-4444","fullName":"Lorenzo  
Phil"}  
{"ticketNo":1762320569757,"contactPhone":"364-610-4455","fullName":"Lucinda  
Beckman"}
```

User data application table

The following code creates the user data application table - users table.

```
CREATE TABLE users ( id INTEGER,  
                      firstName STRING,  
                      lastName STRING,  
                      otherNames ARRAY(RECORD(first STRING, last STRING)),  
                      age INTEGER,  
                      income INTEGER,
```

```

        address JSON,
        connections ARRAY(INTEGER),
        email STRING,
        communityService STRING,
PRIMARY KEY (id) )

```

The following code populates users tables with sample rows.

```

INSERT INTO users VALUES (

    10,
    "John",
    "Smith",
    [ {"first" : "Johny", "last" : "Good"}, {"first" : "Johny2", "last" :
"Brave"}, {"first" : "Johny3", "last" : "Kind"}, {"first" : "Johny4", "last" :
"Humble"} ],
    22,
    45000,
    {
        "street" : "Pacific Ave",
        "number" : 101,
        "city" : "Santa Cruz",
        "state" : "CA",
        "zip" : 95008,
        "phones" : [
            { "area" : 408, "number" : 4538955, "kind" : "work" },
            { "area" : 831, "number" : 7533341, "kind" : "home" },
            { "area" : 831, "number" : 7533382, "kind" : "mobile" }
        ]
    },
    [ 30, 55, 43 ],
    "john.smith@reachmail.com",
    "*****"
)

```

```

INSERT INTO users VALUES (

    20,
    "Jane",
    "Smith",
    [ {"first" : "Jane", "last" : "BeGood"} ],
    22,
    55000,
    {
        "street" : "Atlantic Ave",
        "number" : 201,
        "city" : "San Jose",
        "state" : "CA",
        "zip" : 95005,
        "phones" : [
            { "area" : 608, "number" : 6538955, "kind" : "work" },
            { "area" : 931, "number" : 9533341, "kind" : "home" },
            { "area" : 931, "number" : 9533382, "kind" : "mobile" }
        ]
    },
    [

```

```

        [ 40, 75, 63 ],
        "jane.smith201@reachmail.com",
        "*****"
    )

INSERT INTO users VALUES (

    30,
    "Adam",
    "Smith",
    [ {"first" : "Adam", "last" : "BeGood"} ],
    45,
    75000,
    {
        "street" : "Indian Ave",
        "number" : 301,
        "city" : "Houston",
        "state" : "TX",
        "zip" : 95075,
        "phones" : [
            { "area" : 618, "number" : 6618955, "kind" : "work" },
            { "area" : 951, "number" : 9613341, "kind" : "home" },
            { "area" : 981, "number" : 9613382, "kind" : "mobile" }
        ]
    },
    [ 60, 45, 73 ],
    "adam.smith201reachmail.com",
    "****"
)

INSERT INTO users VALUES (

    40,
    "Joanna",
    "Smith",
    [ {"first" : "Joanna", "last" : "Smart"} ],
    NULL,
    75000,
    {
        "street" : "Tex Ave",
        "number" : 401,
        "city" : "Houston",
        "state" : "TX",
        "zip" : 95085,
        "phones" : [
            { "area" : NULL, "number" : 1618955, "kind" : "work" },
            { "area" : 451, "number" : 4613341, "kind" : "home" },
            { "area" : 481, "number" : 4613382, "kind" : "mobile" }
        ]
    },
    [ 70, 30, 40 ],
    NULL,
    "***"
)

INSERT INTO users VALUES (

```

```

50,
"Aubrey",
"Winslet",
[ {"first" : "Aubrey", "last" : "Regal"} ],
NULL,
15000,
{
  "street" : "Tex Ave",
  "number" : 501,
  "city" : "Houston",
  "state" : "TX",
  "zip" : 95085,
  "phones" : [
    { "area" : 723, "number" : 7618955, "kind" : "work" },
    { "area" : 751, "number" : 7613341, "kind" : "home" },
    { "area" : 781, "number" : 7613382, "kind" : "mobile" }
  ]
},
[ 50, 20, 40 ],
"reachaubrey@somemail.co.us",
"*****"
)

```

```
INSERT INTO users VALUES (
```

```

60,
"Jimmy",
"Law",
[ {"first" : "Jimmy", "last" : "Smart"} ],
NULL,
25000,
{
  "street" : "Maine",
  "number" : 501,
  "city" : "San Jose",
  "state" : "TX",
  "zip" : 95085,
  "phones" : [
    { "area" : 223, "number" : 2618955, "kind" : "work" },
    { "area" : 251, "number" : 2613341, "kind" : "home" },
    { "area" : 281, "number" : 2613382, "kind" : "mobile" }
  ]
},
[ 50, 20, 40 ],
"reachjimlaw@gotmail.co.us",
"*****"
)

```

```
INSERT INTO users VALUES (
```

```

70,
"Dierdre",
"Amador",
[ {"first" : "Dierdre", "last" : "Kind"} ],
NULL,

```

```

25000,
{
  "street" : "Maine\t(Suburb)",
  "number" : 701,
  "city" : "San Jose",
  "state" : "TX\u001B",
  "zip" : 95085,
  "phones" : [
    { "area" : 223, "number" : 6718955, "kind" : "work" },
    { "area" : 251, "number" : 6213341, "kind" : "home" },
    { "area" : 281, "number" : 6213382, "kind" : "mobile" }
  ]
},
[ 10, 60, 40 ],
NULL,
"****"
)

```

Exists Operator

Syntax

`exists_expression ::= EXISTS concatenate\_expression`

Semantics

The exists operator checks whether the sequence returned by its input expression is empty or not, and returns false or true, respectively. A special case is when the input expression returns NULL. In this case, EXISTS will also return NULL, unless it is known that the input expression will always return at least one item, in which case EXISTS returns true.

Example 6-84 Exists Operator

Find all the users who do not have a zip code in their addresses.

```

SELECT id FROM users u
WHERE NOT EXISTS u.address.zip

```

Notice that the above query does not select users whose zip code has the json null value. The following query includes those users as well.

```

SELECT id FROM users u
WHERE NOT EXISTS u.address.zip OR u.address.zip = null

```

What if the Users table contains a row R whose address column is NULL? In general, SQL for Oracle NoSQL Database interprets NULL as an unknown value, rather than an absent value. So, in row R, the address is unknown, and as a result, we don't know what its zip code is or if it even has a zip code. In this case, the expression `u.address.zip` will return NULL on R and `exists u.address.zip` will also return NULL, which implies that row R will not be selected by the above queries. On the other hand, row R will be selected by the following query. In this case, we know that every row does have an address, even though the address may be unknown

(i.e., NULL) in some rows. So, even though the expression `u.address` returns NULL, `exists u.address` return true.

```
SELECT id FROM users u
WHERE EXISTS u.address
```

Is-Of-Type Operator

Syntax

```
is_of_type_expression ::=
  add\_expression IS [NOT] OF [TYPE]
  "(" [ONLY] sequence\_type ([ONLY] sequence_type)* ")"
```

Semantics

The is-of-type operator checks the sequence type of its input sequence against one or more target sequence types. If the number N of the target types is greater than one, the expression is equivalent to OR-ing N is-of-type expressions, each having one target type. So, for the remainder of this section, we will assume that only one target type is specified.

The is-type-of operator will return true if both of the following conditions are true:

1. the cardinality of the input sequence matches the quantifier of the target type. Specifically,
 - a. if the quantifier is `*` the sequence may have any number of items,
 - b. if the quantifier is `+` the input sequence must have at least one item,
 - c. if the quantifier is `?` The input sequence must have at most one item, and
 - d. if there is no quantifier, the input sequence must have exactly one item.
2. all the items in the input sequence are instances of the target item-type (`type_def`), i.e. the type of each input item must be a subtype of the target item-type. For the purposes of this check, a NULL is not considered to be an instance of any type.

If condition (1) is satisfied and the input sequence contains a NULL, the result of the is-type-of operator will be NULL. In all other cases, the result is false.

Example 6-85 Is-Of-Type Operator

Find all the users whose address information has been stored as a single, flat string.

```
SELECT id
FROM users u
WHERE u.address IS OF TYPE (STRING)
```

Concatenation Operator

Syntax

```
concatenation_operator ::= "||"

concatenate_expression ::= add\_expression ("||" add_expression)*
```


Semantics

The concatenation operator returns the character string made by joining the operands in the given order. The operands can be of any* type. For more details, see the [concat Function](#) section.

Note

According to the operator precedence, the || operator is immediately after +, - (as binary operators).

Example 6-86 Concatenation Operator

This example joins id, firstname, and lastname into a single string and provides the output. Notice that id, which is an integer type, also gets concatenated with the string values.

```
SELECT id || firstname || lastname AS name FROM users
```

Output

```
+-----+
|   name   |
+-----+
| 10JohnSmith |
| 30PeterPaul |
| 20MaryAnn   |
+-----+
```

Arithmetic Expressions

Syntax

```
add_expression ::= multiply_expression (("+"|"") multiply_expression)*
```

```
multiply_expression ::= unary_expression (("*"|"/"|"div") unary_expression)*
```

```
unary_expression ::= path\_expression | (("+"|"") unary_expression)
```

Semantics

Oracle NoSQL Database supports the following arithmetic operations: +, -, *, / and div. Each operand to these operators must produce at most one numeric item. If any operand returns the empty sequence or NULL, the result of the arithmetic operation is also empty or NULL, respectively. Otherwise, the operator returns a single numeric item, which is computed as follows:

- If any operand returns a Number item, the item returned by the other operand is cast to a Number value (if not a Number already) and the result is a Number item that is computed using java's arithmetic on BigDecimal, otherwise,

- If any operand returns a double item, the item returned by the other operand is cast to a double value (if not a double already) and the result is a double item that is computed using java's arithmetic on doubles, otherwise,
- If any operand returns a float item, the item returned by the other operand is cast to a float value (if not a float already) and the result is a float item that is computed using java's arithmetic on floats, otherwise,
- Except for the div operator, if any operand returns a long item, the item returned by the other operand is cast to a long value (if not a long already) and the result is a long item that is computed using java's arithmetic on longs.
- Except for the div operator, if all operands return integer items, the result is an integer item that is computed using java's arithmetic on ints.
- The div operator performs floating-point division, even if both its operands are longs and/or integers. In this case, div returns a double.

Oracle NoSQL Database supports the unary + and – operators as well. The unary + is a no-op, and the unary – changes the sign of its numeric argument.

To follow along with the examples, create the user data and airline baggage tracking application table, and insert the data as described in the [Tables Used in the Examples](#) section.

Example 6-87 Arithmetic Expression

For each user show their id and the difference between their actual income and an income that is computed as a base income plus an age-proportional amount.

```
DECLARE
$baseIncome INTEGER;
$ageMultiplier DOUBLE;
SELECT id,
income - ($baseIncome + age * $ageMultiplier) AS adjustment
FROM Users
```

Example 6-88 Fetch the name, number of transits, and calculate the transit time for the passenger with ticket number 1762320369957 from the airline baggage tracking application data

```
SELECT bag.fullname AS NAME, $t-1 AS HOPS,
CASE
WHEN $t-1 = 2
THEN
get_duration(timestamp_diff($bagInfo.flightLegs[2].flightDate, $bagInfo.flight
Legs[1].estimatedArrival) +
timestamp_diff($bagInfo.flightLegs[1].flightDate, $bagInfo.flightLegs[0].estim
atedArrival))
WHEN $t-1 = 1
THEN
get_duration(timestamp_diff($bagInfo.flightLegs[1].flightDate, $bagInfo.flight
Legs[0].estimatedArrival))
ELSE "Direct flight"
END AS TRANSITHRS
FROM BaggageInfo bag, bag.bagInfo[0] AS $bagInfo,
size($bagInfo.flightLegs) $t
WHERE bag.ticketNo = 1762320369957
```

Explanation: In the `BaggageInfo` table, the JSON field `bagInfo` stores the checked baggage tracking information for each passenger. Passengers can have multiple transits in their air travel. The `flightLegs` array contains the information on each airport in the passenger's travel itinerary. In each element of the `flightLegs` array, the `flightDate` field holds the scheduled departure time from the source airport and the `estimatedArrival` field holds the estimated arrival time at the destination airport. You use the [size](#) function to calculate the size of the `flightLegs` array, which is the number of travel legs in the passenger's air travel. However, this count includes even the last airport, which is the destination and not the transit airport. To calculate the number of transit airports (hops) per passenger, you use the arithmetic operator '-' to subtract the last airport from the count. Notice that you can assign the size calculation to an internal variable `$t1` to facilitate the usage throughout the query. For more details on variables, see [Variable Declaration](#).

To calculate the time spent by a passenger in each transit, you find the duration between the `flightDate` of the next immediate destination airport and the `estimatedArrival` of the current airport. You use the [timestamp_diff](#) function to calculate the time difference. Repeat this for all the transits until the `flightDate` field belongs to the final destination airport. Use the arithmetic operator '+' to add the individual transit times to arrive at the total transit time. Use the [get_duration](#) function to convert the result from milliseconds to a duration string. Since the number of hops is not fixed for all the passengers, you use the [CASE expression](#) with a condition each for the number of hops. Depending on the number of hops for a given passenger, the corresponding condition applies, and the total transit time is calculated and displayed in the result. Here, only two hops are considered for ease of calculation.

The passengers may have more than one checked bag, in which case there will be more than one element in the `bagInfo` array. You must consider only the first element of the `bagInfo` array, that is, `bagInfo[0]` during calculation to avoid duplication of results.

Output:

```
{"NAME":"Lorenzo Phil","TRANSITHRS":"6 hours 8 minutes"}
```

Primary Expressions

Syntax

```
primary_expression ::=  
    parenthesized\_expression |  
    constant\_expression |  
    column\_reference |  
    variable\_reference |  
    array\_constructor |  
    map\_constructor |  
    case\_expression |  
    cast\_expression |  
    extract\_expression |  
    function\_call |  
    count\_star |  
    transform\_expression
```

The following subsections describe each of the primary expressions listed in the above grammar rule, except from `count_star`, which is the `count(*)` aggregate function defined in the [Using Aggregate Functions](#) section.

To follow along with the examples, create tables and insert the data as described in the [Tables Used in the Examples](#) section.

Parenthesized Expressions

Syntax

parenthesized_expression ::= "(" [expression](#) ")"

Semantics

Parenthesized expressions are used primarily to alter the default precedence among operators. They are also used as a syntactic aid to mix expressions in ways that would otherwise cause syntactic ambiguities. An example of the later usage is in the definition of the `field_step` parse rule in the [Field Step Expressions](#) section.

Example 6-89 Parenthesized Expression

Select the id and the last name for users whose age is less or equal to 30 and either their age is greater than 20 or their income is greater than 100K.

```
SELECT id, lastName
FROM users
WHERE (income > 100000 OR 20 < age) AND age <= 30
```

Constant Expressions

Syntax

constant_expression ::= number | string | TRUE | FALSE | NULL

number ::= [MINUS] (FLOAT_CONSTANT | INT_CONSTANT | NUMBER_CONSTANT)

string ::= STRING_CONSTANT | DSTRING_CONSTANT

Semantics

The syntax for INT_CONSTANT, FLOAT_CONSTANT, NUMBER_CONSTANT, STRING_CONSTANT, and DSTRING_CONSTANT was given in the [Identifiers](#) section.

In the current version, a query can contain the following constants (a.k.a. literals):

String

String literals are sequences of unicode characters enclosed in double or single quotes. String literals are translated into String items. Notice that any escape sequences appearing in a string literal will be converted to their corresponding character inside the corresponding String item.

Integer

Integer literals are sequences of one or more digits. Integer literals are translated into Integer items, if their value fits in 4 bytes, into Long items, if they fit in 8 bytes, otherwise to Number items.

Floating point

Floating point literals represent real numbers using decimal notation and/or exponent. Floating-point literals are translated into Double items, if possible, otherwise to Number items.

Number

Number literals are integer or floating-point literals followed by the 'n' or 'N' character. Number literals are always translated into Number items.

TRUE / FALSE

The TRUE and FALSE literals are translated to the boolean true and false items, respectively.

NULL

The NULL literal is translated to the json null item.

Example 6-90 From the airline baggage tracking application data, fetch the list of passengers whose bags were delivered to JTR airport

```
SELECT fullname, bag.bagInfo.id, bag.bagInfo.tagNum
FROM BaggageInfo bag
WHERE bag.bagInfo.lastSeenStation =any "JTR"
```

Explanation: In the BaggageInfo table, the lastSeenStation field contains information on the current location of the passenger's checked luggage. In this query, you fetch the ID and tag numbers of bags that were delivered to the JTR airport. You use the string literal "JTR" as a constant expression in the WHERE clause. Here, the bagInfo is an array and bag.bagInfo.lastSeenStation is a sequence. You want to retrieve any element in the sequence whose lastSeeStation was JTR. As the [Value Comparison Operators](#) cannot operate on sequences of more than one item, you use the [sequence comparison operator](#) any in addition to the value comparison operator '=' to compare the lastSeeStation fields.

If a passenger carries more than one bag, the ID and tag numbers of both bags are listed in the result.

Output:

```
{"fullname":"Gerard Greene","id":"79039899152842","tagNum":"1765780626568"}
{"fullname":"Michelle Payne","id":"79039899186259","tagNum":"17657806247861"}
{"fullname":"Lorenzo Phil","id":["79039899187755","79039899197755"],"tagNum":
["17657806240001","17657806340001"]}
{"fullname":"Lucinda Beckman","id":"79039899187755","tagNum":"17657806240001"}
```

Column References

Syntax

column_reference ::= [id](#) ["." [id](#)]

Semantics

A column-reference expression returns the item stored in the specified column within the context row (the row that an SELECT expression is currently working on). Syntactically, a column-reference expression consists of one identifier, or 2 identifiers separated by a dot. If there are 2 ids, the first is considered to be a table alias and the second the name of a column in that table. We call this form a qualified column name. A single id is resolved to the name of a column in some of the tables referenced inside the FROM clause. However, in this case there

must not be more than one tables that participate in the query and have a column with this name. We call this form an unqualified column name.

Example 6-91 Fetch the bag arrival date for a passenger from the airline baggage tracking application data

```
SELECT fullname, bag.bagInfo[0].bagArrivalDate
FROM BaggageInfo bag
WHERE ticketNo=1762392135540
```

Explanation: In the BaggageInfo table, the bagArrivalDate field contains the date on which the checked baggage is delivered to the destination airport. To fetch the bag arrival date, you use a qualified column name in your query, that is, bag.bagInfo[0].bagArrivalDate. For more details, see [Variable Declaration](#).

Here, the first reference bag is a table alias to the BaggageInfo table. The second reference bagInfo[0] is to the first element of the bagInfo JSON array column. The third reference is to the actual field bagArrivalDate that is nested inside the bagInfo JSON column.

You use only the first element of the bagInfo JSON array to avoid duplication of results for passengers carrying more than one bag.

Output:

```
{"fullname": "Adelaide
Willard", "bagArrivalDate": "2019-02-15T21:21:00Z"}
```

Variable References

Syntax

```
variable_reference ::= "$" [id]
```

Semantics

A variable-reference expression returns the item that the specified variable is currently bound to. Syntactically, a variable-reference expression is just the name of the variable.

Oracle NoSQL Database allows the declaration of both external and internal variables. For more details on declaring the variables and their scope, see [Variable Declaration](#).

The table aliases are essentially internal variables ranging over the rows of the specified table. The other clauses of the SELECT expression operate on the rows produced by the FROM clause, processing one row at a time. The row currently being processed is called the context row. The columns of the context row can be referenced in expressions either directly by their names or by the table alias followed by a dot character and the column name. If the table alias starts with a dollar sign (\$), then it serves as an internal variable declaration for a variable whose name is the alias. This internal variable is bound to the context row as a whole and can be referenced within sub-expressions of the SELECT expression. It can also be passed as an argument to a function to fetch the result for the context row.

Example 6-92 Fetch the duration between the baggage boarding and baggage collection from the destination airport for a passenger from the airline baggage tracking application data

```

SELECT $flightLeg.flightDate AS BAGGAGE_CHECK, $bagInfo.bagArrivalDate AS
BAGGAGE_COLLECT,
$t1 AS BAGGAGE_DURATION
FROM BaggageInfo bag, bag.bagInfo[0] AS $bagInfo, $bagInfo.flightLegs[0]
AS $flightLeg,
get_duration(timestamp_diff($bagInfo.bagArrivalDate, $flightLeg.flightDate)) $
t1
WHERE ticketNo=1762320369957

```

Explanation: In this example, you use table alias and internal variables. You can determine the duration between the checked time and the baggage collection time at the destination airport by calculating the time difference between `bagArrivalDate` and `flightDate[0]` for each passenger. You use the `timestamp_diff` function to find the time difference and the `get_duration` function to convert the result from milliseconds to a duration string. For more details, see [timestamp_diff and get_duration functions](#).

The table name `BaggageInfo` in the query is followed by a table alias `bag`. You use the variables `$bagInfo` and `$flightLeg` to reference the columns in the context row (the row that is currently being processed). Notice that `bagInfo` is a column of type array in the `BaggageInfo` table, whereas `$bagInfo` is an internal variable used to hold the first element of the `bagInfo` array for the context row. `$t1` is an internal variable whose value is assigned when the query is executed.

Note

Here, the `timestamp_diff` function automatically casts the `bagArrivalDate` and `flightDate` from String to a timestamp data type during the query run time.

Output:

```

{"BAGGAGE_CHECK":"2019-03-12T03:00:00Z","BAGGAGE_COLLECT":"2019-03-12T15:05:00
Z","BAGGAGE_DURATION":"12
hours 5 minutes"}

```

Array and Map Constructors

Syntax

```

array_constructor ::=
    "[" expression ("," expression)* "]"

map_constructor ::=
    ("{" expression ":" expression
    ("," expression ":" expression)* "}") |
    ("{" "}")

```

Semantics

An array constructor constructs a new array out of the items returned by the expressions inside the square brackets. These expressions are computed left to right, and the produced items are appended to the array. Any NULLs produced by the input expressions are skipped (arrays cannot contain NULLs).

Similarly, a map constructor constructs a new map out of the items returned by the expressions inside the curly brackets. These expressions come in pairs: each pair computes one field. The first expression in a pair must return at most one string, which serves as the field's name and the second returns the associated field value. If a value expression returns more than one items, an array is implicitly constructed to store the items, and that array becomes the field value. If either a field name or a field value expression returns the empty sequence, no field is constructed. If the computed name or value for a field is NULL the field is skipped (maps cannot contain NULLs).

The type of the constructed arrays or maps is determined during query compilation, based on the types of the input expressions and the usage of the constructor expression. Specifically, if a constructed array or map may be inserted in another constructed array or map and this "parent" array/map has type ARRAY(JSON) or MAP(JSON), then the "child" array/map will also have type ARRAY(JSON) or MAP(JSON). This is to enforce the restriction that "typed" data are not allowed inside JSON data (see [Data Type Hierarchy](#)).

Example 6-93 Array and Map Constructor

For each user create a map with 3 fields recording the user's last name, their phone information, and the expense categories in which more than \$5000 was spent.

```
SELECT
{
  "last_name" : u.lastName,
  "phones" : u.address.phones,
  "high_expenses" : [ u.expenses.keys($value > 5000) ]
}
FROM users u
```

Notice that the use of an explicit array for the "high_expenses" field guarantees that the field will exist in all the constructed maps, even if the path inside the array constructor returns empty. Notice also that although it is known at compile time that all elements of the constructed arrays will be strings, the arrays are constructed with type ARRAY(JSON) (instead of ARRAY(STRING)), because they are inserted into a JSON map.

Case Expressions

Syntax

```
case_expression ::= CASE
  WHEN expression THEN expression
  (WHEN expression THEN expression)*
  [ELSE expression]
END
```


Semantics

The searched CASE expression is similar to the if-then-else statements of traditional programming languages. It consists of a number of WHEN-THEN pairs, followed by an optional ELSE clause at the end. Each WHEN expression is a condition, i.e., it must return `BOOLEAN`. The THEN expressions as well as the ELSE expression may return any sequence of items. The CASE expression is evaluated by first evaluating the WHEN expressions from top to bottom until the first one that returns true. If it is the i-th WHEN expression that returns true, then the i-th THEN expression is evaluated and its result is the result of the whole CASE expression. If no WHEN expression returns true, then if there is an ELSE, its expression is evaluated and its result is the result of the whole CASE expression; Otherwise, the result of the CASE expression is the empty sequence.

Example 6-94 For each user create a map with 3 fields recording the user's last name, their phone information, and the expense categories in which more than \$5000 was spent

```
SELECT
{
  "last_name" : u.lastName,
  "phones" : CASE
    WHEN exists u.address.phones
    THEN u.address.phones
    ELSE "Phone info absent or not at the expected place"
  END,
  "high_expenses" : [ u.expenses.keys($value > 5000) ]
}
FROM users u
```

The query is very similar to the one from array and map constructor. The only difference is in the use of a case expression to compute the value of the phones field. This guarantees that the phones field will always be present, even if the path expression `u.address.phones` return empty or NULL. Notice that wrapping the path expression with an explicit array constructor (as we did for the `high_expenses` field) would not be a good solution here, because in most cases `u.address.phones` will return an array, and we don't want to have construct an extra array containing just another array.

JSON collection table:

The following example applies CASE expression on a JSON collection table. Consider a sample row from the JSON collection table created for a [shopping application](#):

```
{"contactPhone":"1917113999","address":{"city":"San
Jose","number":501,"state":"San
Francisco","street":"Maine","zip":95095},"cart":
[{"item":"wallet","priceperunit":950,"quantity":2},{
"item":"wall art","priceperunit":9500,"quantity":1}],
"firstName":"Sharon","gender":"F","lastName":"Willard",
"notify":"yes","wishlist":
[{"item":"Tshirt","priceperunit":500},{
"item":"Jenga","priceperunit":850}]}
```

Example 6-95 Display promotional messages to shoppers from San Jose who have wallet or handbag items in their carts

```
SELECT concat("Hi ",s.firstName) AS Message,
CASE
```

```

        WHEN s.cart.item =any "wallet"
        THEN "The prices on Wallets have dropped"
        WHEN s.cart.item =any "handbag"
        THEN "The prices on handbags have dropped"
        ELSE "Exciting offers on wallets and handbags"
    END AS Offer
FROM storeAcct s

WHERE s.address.city =any "San Jose"

```

Explanation: You can use CASE statement to display a promotional message to the shoppers regarding the reduction in the prices if the shoppers have the items wallet or handbag in their cart. As the offers are only for shoppers from San Jose, you specify the city in the WHERE clause.

Output:

```

{"Message":"Hi Sharon","Offer":"The prices on Wallets have
dropped"}

```

Cast Expression

Syntax

```
cast_expression ::= CAST "(" expression AS sequence\_type ")"
```

Semantics

The cast expression creates, if possible, new items of a given target type from the items of its input sequence. Specifically, a cast expression is evaluated as follows:

A cardinality check is performed first:

1. if the quantifier of the target type is * the sequence may have any number of items,
2. if the quantifier is + the input sequence must have at least one item,
3. if the quantifier is ? the input sequence must have at most one item, and
4. if there is no quantifier, the input sequence must have exactly one item.

If the cardinality of the input sequence does not match the quantifier of the target type, an error is raised. Then, each input item is cast to the target item type according to the following (recursive) rules.

- If the type of the input item is equal to the target item type, the cast is a no-op: the input item itself is returned.
- If the target type is a wildcard type other than JSON and the type of the input item is a subtype of the wildcard type, the cast is a no-op; Otherwise an error is raised.
- If the target type is JSON, then (a) an error is raised if the input item is has a non-json atomic type, else (b) if the input item has a type that is a json atomic type or ARRAY(JSON) or MAP(JSON), the cast is a no-op , else (c) if the input item is a non-json array, a new array of type ARRAY(JSON) is constructed, each element of the input array is cast to JSON, and the resulting item is appended into the new json array, else (d) if the input item is a non-json map, a new map of type MAP(JSON) is constructed, each field value of the input map is cast to JSON, and resulting item together with the associated

field name are inserted into the new json map, else (e) if the input item is a record, it is cast to a map of type MAP(JSON) as described below.

- If the target type is an array type, an error is raised if the input item is not an array. Otherwise, a new array is created, whose type is the target type, each element in the input array is cast to the element type of the target array, and the resulting item is appended into the new array.
- If the target type is a map type, an error is raised if the input item is not a map or a record. Otherwise, a new map is created, whose type is the target type, each field value in the input map/record is cast to the value type of the target map, and the resulting field value together with the associated field name are inserted to the new map.
- If the target type is a record type, an error is raised if the input item is not a record or a map. Otherwise, a new record is created, whose type is the target type. If the input item is a record, its type must have the same fields and in the same order as the target type. In this case, each field value in the input record is cast to the value type of the corresponding field in the target type and the resulting field value together with the associated field name are added to the new record. If the input item is a map, then for each map field, if the field name exists in the target type, the associated field value is cast to the value type of the corresponding field in the target type and the resulting field value together with the associated field name are added to the new record. Any fields in the new record whose names do not appear in the input map have their associated field values set to their default values.
- If the target type is string, the input item may be of any type. In other words, every item can be cast to a string. For complex items their “string value” is a json-text representation of their value. For timestamps, their string value is in UTC and has the format "YYYY-MM-dd['T'HH:mm:ss]". For binary items, their string value is a base64 encoding of their bytes.
- If the target type is an atomic type other than string, the input item must also be atomic. Among atomic items and types the following casts are allowed:
 - Every numeric item can be cast to every other numeric type. The cast is done as in Java.
 - Integers and longs can be cast to timestamps. The input value is interpreted as the number of milliseconds since January 1, 1970, 00:00:00 GMT.
 - String items may be castable to all other atomic types. Whether the cast succeeds or not depends on whether the actual string value can be parsed into a value that belongs to the domain of the target type.
 - Timestamp items are castable to all the timestamp types. If the target type has a smaller precision than the input item, the resulting timestamp is the one closest to the input timestamp in the target precision. For example, consider the following 2 timestamps with precision 3: 2016-11-01T10:00:00.236 and 2016-11-01T10:00:00.267. The result of casting these timestamps to precision 1 is: 2016-11-01T10:00:00.2 and 2016-11-01T10:00:00.3, respectively. If the target type has a larger precision than the input item, additional zeros are appended to the resulting timestamp value.

Example 6-96 Cast Expression

Select the last name of users who moved to their current address in 2015 or later.

```
SELECT u.lastName FROM Users u
WHERE
CAST (u.address.startDate AS Timestamp(0)) >=
CAST ("2015-01-01T00:00:00" AS Timestamp(0))
```

Since there is no literal for Timestamp values, to create such a value a string has to cast to a Timestamp type.

Extract Expressions

Syntax

```
extract_expression ::= EXTRACT "(" id FROM expression ")"
```

Semantics

The extract expression extract a component from a timestamp. Specifically, the expression after the FROM keyword must return at most one timestamp or NULL. If the result of this expression is NULL or empty, the results of EXTRACT is also NULL or empty, respectively. Otherwise, the component specified by the id is returned. This id must be one of the following keywords:

YEAR

Returns the year for the timestamp, in the range -6383 ~ 9999.

MONTH

Returns the month for the timestamp, in the range 1 ~ 12.

DAY

Returns the day of month for the timestamp, in the range 1 ~ 31.

HOURL

Returns the hour of day for the timestamp, in the range 0 ~ 23.

MINUTE

Returns the minute for the timestamp, in the range 0 ~ 59.

SECOND

Returns the second for the timestamp, in the range 0 ~ 59.

MILLISECOND

Returns the fractional second in millisecond for the timestamp, in the range 0 ~ 999.

MICROSECOND

Returns the fractional second in microsecond for the timestamp, in the range 0 ~ 999999.

NANOSECOND

Returns the fractional second in nanosecond for the timestamp, in the range 0 ~ 999999999.

WEEK

Returns the week number within the year where a week starts on Sunday and the first week has a minimum of 1 day in this year, in the range 1 ~ 54.

ISOWEEK

Returns the week number within the year based on ISO-8601, where a week starts on Monday and the first week has a minimum of 4 days in this year, in range 0 ~ 53.

There are specific built-in functions to extract each of the above components from a time stamp. For example, EXTRACT(YEAR from expr) is equivalent to year(expr). These and other built-in functions are described in [Built-in Functions](#).

Example 6-97 Create a message that displays the bag arrival date, time, and destination airport details for a passenger from the airline baggage tracking application data

```
SELECT concat("Hi ", bag.fullname, ", Please collect your bags from  
", $station, " airport on ", $t2, "/", $t3, "/", $t1, " at ", $t4, ":", $t5, " Hrs") AS  
MESSAGE  
FROM BaggageInfo bag, bag.bagInfo[0].bagArrivalDate AS $bagDate,  
bag.bagInfo[0].lastSeenStation AS $station,  
EXTRACT(YEAR from (CAST($bagDate AS Timestamp(0)))) $t1,  
EXTRACT(MONTH from (CAST($bagDate AS Timestamp(0)))) $t2,  
EXTRACT(DAY from (CAST($bagDate AS Timestamp(0)))) $t3,  
EXTRACT(HOUR from (CAST($bagDate AS Timestamp(0)))) $t4,  
EXTRACT(MINUTE from (CAST($bagDate AS Timestamp(0)))) $t5  
WHERE ticketNo = 1762344493810
```

Explanation: The `bagArrivalDate` field in the `BaggageInfo` table stores the date on which the checked bags are scheduled to arrive at the destination airport. You can use the `EXTRACT` expression to extract the date (day, month, and year) and time (hour and minutes) from the `bagArrivalDate` field for the specified passenger. The date is stored as a string in the JSON data in the `bagInfo` field. You first use the [CAST expression](#) to convert the `bagArrivalDate` to a timestamp and then extract the date, month, year, and so forth from the timestamp.

To display as a message, you use the [concat](#) function.

Output:

```
{"MESSAGE": "Hi Adam Phillips, Please collect your bags from MEL airport on  
2/1/2019 at  
16:13 Hrs"}
```

Function Calls

Syntax

```
function_call ::= id "(" [expression ("," expression)*] ")"
```

Semantics

Function-call expressions are used to invoke functions, which in the current version can be built-in (system) functions only. Syntactically, a function call starts with an id which identifies the function to call by name, followed by a parenthesized list of zero or more argument expressions separated by a comma.

Each function has a signature, which specifies the sequence type of its result and a sequence type for each of its parameters. Evaluation of a function-call expression starts with the evaluation of each of its arguments. The result of each argument expression must be a subtype of the corresponding parameter type, or otherwise, it must be promotable to the parameter type. In the latter case, the argument value will actually be cast to the expected type. Finally, after type checking and any necessary promotions are done, the function's implementation is invoked with the possibly promoted argument values.

The following type promotions are currently supported:

- INTEGER is promotable to FLOAT or DOUBLE.

- LONG is promotable to FLOAT or DOUBLE.
- STRING is promotable to ENUM, but the cast will succeed only if the ENUM type contains a token whose string value is the same as the input string.

See the [Built-in Functions](#) topic for the list of the supported functions in Oracle NoSQL Database.

Example 6-98 Fetch the full names of passengers who have the airport code MEL in their route from the airline baggage tracking application data

```
SELECT fullname
FROM baggageInfo bag
WHERE EXISTS bag.bagInfo[contains($element.routing, "MEL")]
```

Explanation: In the BaggageInfo table, the routing field contains the codes of the airports through which the checked bag transits before reaching the destination airport. In the query above, you use a function call to the [contains](#) function. The contains function is one of the built-in functions, which indicates whether or not a search string is present inside the source string. The square brackets in the query iterates over the elements of the bagInfo array. During the iteration, the \$element variable is bound to the current array element, that is, routing. Each iteration computes the expression inside the contains function on the array element. If the expression includes the string "MEL", it returns true and the element is included in the result, otherwise it is skipped. As a result, only the rows that include the MEL airport code in their routing field are displayed in the output.

Output:

```
{"fullname":"Zulema Martindale"}
{"fullname":"Adam Phillips"}
{"fullname":"Joanne Diaz"}
{"fullname":"Omar Harvey"}
{"fullname":"Zina Christenson"}
```

Sequence Transform Expressions

Syntax

```
transform_expression ::= SEQ_TRANSFORM "(" expression ", " expression ")"
```

Semantics

A sequence transform expression transforms a sequence to another sequence. Syntactically it looks like a function whose name is seq_transform. The first argument is an expression that generates the sequence to be transformed (the input sequence) and the second argument is a "mapper" expression that is computed for each item of the input sequence. The result of the seq_transform expression is the concatenation of sequences produced by each evaluation of the mapper expression. The mapper expression can access the current input item via the \$ variable.

Example 6-99 Sequence Transform Expression

As an example, assume a "sales" table with the following data.

```
CREATE TABLE sales (
  id INTEGER,
  sale RECORD (
    acctno INTEGER,
    year INTEGER,
    month INTEGER,
    day INTEGER,
    state STRING,
    city STRING,
    storeid INTEGER,
    prodcats STRING,
    items ARRAY(
      RECORD (
        prod STRING,
        qty INTEGER,
        price INTEGER
      )
    )
  ),
  PRIMARY KEY (id)
)

INSERT INTO sales VALUES (
  1,
  {
    "acctno" : 349,
    "year" : 2000,
    "month" : 10,
    "day" : 23,
    "state" : "CA",
    "city" : "San Jose",
    "storeid" : 76,
    "prodcats" : "vegies",
    "items" :[
      { "prod" : "tomatoes", "qty" : 3, "price" : 10.0 },
      { "prod" : "carrots", "qty" : 1, "price" : 5.0 },
      { "prod" : "pepers", "qty" : 1, "price" : 15.0 }
    ]
  }
)
```

Assume there is the following index on sales:

```
CREATE INDEX idv1 ON sales (
  sale.acctno, sale.year, sale.prodcats)
```

Then we can write the following query, which returns the total sales per account number and year:

```
SELECT t.sale.acctno,
       t.sale.year,
       sum(seq_transform(t.sale.items[], $.price * $.qty)) AS sales
FROM sales t
GROUP BY t.sale.acctno, t.sale.year
```

Using sequence transform expression for JSON documents :

You can use the sequence transform expression for transforming JSON documents stored in table rows. In such cases you often use multiple sequence transform expressions nested inside each other. Here the mapper expression of an inner sequence transform may need to access the current item of an outer sequence transform. To allow this, each sequence transform expression 'S' declares a variable with name \$sqN, where N is the level of nesting of the expression 'S' within the outer sequence transform expressions. \$sqN is basically a synonym for \$, that is, it is bound to the items returned by the input expression 'S'. However, \$sqN can be accessed by other sequence transform expressions that may be nested inside the expression 'S'.

You can illustrate this with the following example:

Example 6-100 For each ticketNo, fetch a flat array containing all the actions performed on the luggage of that ticketNo from the airline baggage tracking application.

In this query, you fetch the at and action fields for each action. Also display the flightNo and the tagNum with each action. You can use the sequence transform expression to get a flat array of actions.

```
SELECT
  seq_transform(
    l.bagInfo[],
    seq_transform(
      $sq1.flightLegs[],
      seq_transform(
        $sq2.actions[],
        {
          "at" : $sq3.at,
          "action" : $sq3.action,
          "flightNo" : $sq2.flightNo,
          "tagNum" : $sq1.tagNum
        }
      )
    )
  ) AS actions
FROM baggageInfo l
```

Output:

```
{
  "actions" : [
    { "at": "SF0", "action": "TagScan", "flightNo": "BM604",
      "tagNum": 17657806243915 },
    { "at": "SF0", "action": "onload", "flightNo": "BM604",
      "tagNum": 17657806243915 },
```



```

        {"at":"AMS", "action":"offload", "flightNo":"BM604",
"tagNum":17657806243915},
        {"at":"AMS", "action":"TagScan", "flightNo":"BM667",
"tagNum":17657806243915},
        {"at":"AMS", "action":"onload", "flightNo":"BM667",
"tagNum":17657806243915},
        {"at":"HER", "action":"offload", "flightNo":"BM667",
"tagNum":17657806243915},
        {"at":"SFO", "action":"TagScan", "flightNo":"BM604",
"tagNum":17657806244523},
        {"at":"SFO", "action":"onload", "flightNo":"BM604",
"tagNum":17657806244523},
        {"at":"AMS", "action":"offload", "flightNo":"BM604",
"tagNum":17657806244523},
    ]
}

```

Unnest Arrays & Maps

Use `unnest_syntax` to flatten rows that include arrays or maps.

Syntax of unnest

Syntactically, unnesting is specified as a list of expressions (with associated variables) in the FROM clause of a SELECT statement. Some or all of these expressions may be enclosed in an UNNEST clause.

```

unnest_expression ::= expression AS VARNAME
unnest_clause    ::= UNNEST((unnest_expression)*)
unnest_syntax    ::= (unnest_expression | unnest_clause)*

```

Semantics

The FROM clause of a SELECT statement may contain a list of expressions (with associated variables). Normally, you will use expressions that unnest arrays or maps and the FROM clause will create a new set of rows out of the values of the unnested arrays/maps. Some or all of these expressions may be enclosed in an UNNEST clause. Semantically, the unnest operator is a no-op, that is, whether an UNNEST clause is used or not does not change the result of the FROM clause, that is `unnest(expr)` is the same as `expr`. The purpose of the UNNEST clause is to act as an optimization hint. Specifically, when there is an index on the arrays/maps that are being unnested, the index may be used by the query only if the query uses the UNNEST clause. Furthermore, to help the query processor in using such an index, the UNNEST clause places some restrictions on the expressions that can appear inside it.

You normally use unnesting when you want to group by a field that is inside an array/map. However the two examples below illustrate the semantics of unnesting expressions. So no GROUP BY clause is used in the queries below.

Using a single unnest expression in the FROM clause (to unnest a single map):

Create a table with two columns , one of data type INTEGER, which is the primary key column and second of JSON data type.

Note

The JSON column used in the example is a map.

```
create table sample_unnest ( samp_id INTEGER, samp_data JSON, PRIMARY
KEY(samp_id))

INSERT INTO sample_unnest VALUES(1,
{ "episodeID" : 20, "lengthMin" : 40, "minWatched" : 40 })

INSERT INTO sample_unnest VALUES(2,
{ "episodeID" : 25, "lengthMin" : 20, "minWatched" : 18 })

select * from sample_unnest
      $s1,unnest($s1.samp_data.values() as $s2) where $s1.samp_id=1
```

The above FROM clause references the table sample_unnest by the variable \$s1 and the values of the map samp_data by the variable \$s2 . Conceptually, the result of the FROM clause is a table temp_tbl with two columns. Column 1 stores all rows of table sample_unnest and Column 2 stores all values of the elements of the samp_data map.

Table 6-4 Records in the temp_tbl table

\$s1	\$s2
"samp_id":1,"samp_data": { "episodeID":20, "lengthMin":40, "minWatched":40 }	20
"samp_id":1,"samp_data": { "episodeID":20, "lengthMin":40, "minWatched":40 }	40
"samp_id":1,"samp_data": { "episodeID":20, "lengthMin":40, "minWatched":40 }	40
"samp_id":2,"samp_data": { "episodeID":25, "lengthMin":20, "minWatched":18 }	25
"samp_id":2,"samp_data": { "episodeID":25, "lengthMin":20, "minWatched":18 }	20
"samp_id":2,"samp_data": { "episodeID":25, "lengthMin":20, "minWatched":18 }	18

To see how the rest of the query is executed, take every row from the above table and apply the WHERE condition. For those rows where the condition evaluates to TRUE (\$s1.samp_id=1), the corresponding row is included in the query result. The condition evaluates to TRUE for the first three rows and to FALSE for the last three rows.

Table 6-5 Query output

\$s1	\$s2
{"samp_id":1, "samp_data": {"episodeID":20,"lengthMin":40,"minWatc hed":40}}	20
{"samp_id":1, "samp_data": {"episodeID":20,"lengthMin":40,"minWatc hed":40}}	40
{"samp_id":1, "samp_data": {"episodeID":20,"lengthMin":40,"minWatc hed":40}}	40

Using two unnest expressions in the FROM clause (to unnest an array and the maps nested under the array)

You can use more than one path expression in an UNNEST clause. Let "M" be the number of from-expressions, then the result table of the FROM clause is computed in M steps, each producing an intermediate table.

 Note

The JSON column used in the example is a array of maps.

```
create table sample_unnest (  
    samp_id INTEGER, samp_data JSON, PRIMARY KEY(samp_id))
```

```
INSERT INTO sample_unnest VALUES(1,  
{  
    "episodes":[  
        {  
            "episodeID" : 10,  
            "lengthMin" : 40,  
            "minWatched" : 25  
        },  
        {  
            "episodeID" : 20,  
            "lengthMin" : 35,  
            "minWatched" : 30  
        }  
    ]  
})
```

```
INSERT INTO sample_unnest VALUES(2,  
{  
    "episodes":[  
        {  
            "episodeID" : 10,  
            "lengthMin" : 40,  
            "minWatched" : 25  
        },  
        {  
            "episodeID" : 20,  
            "lengthMin" : 35,  
            "minWatched" : 30  
        }  
    ]  
})
```

```
        "episodeID" : 30,
        "lengthMin" : 40,
        "minWatched" : 25
    },
    {
        "episodeID" : 40,
        "lengthMin" : 35,
        "minWatched" : 30
    }
]
}
)
```

```
INSERT INTO sample_unnest VALUES(3,
{
    "episodes":[
        {
            "episodeID" : 10,
            "lengthMin" : 40,
            "minWatched" : 25
        },
        {
            "episodeID" : 20,
            "lengthMin" : 35,
            "minWatched" : 30
        }
    ]
}
)
```

Query using two path expressions in the UNNEST clause:

```
SELECT n.samp_id as customer,
       $epi.episodeID,$epi.minWatched AS length,$epiVal AS episode_details
FROM   sample_unnest n,
       unnest(n.samp_data.episodes[] AS $epi, $epi.values() AS $epiVal)
```

Step 1 : Records in the first table temp_tbl1 which is the result of first two expressions in the FROM clause (sample_unnest n, unnest(n.samp_data.episodes[] AS \$epi)

Table 6-6 Records in temp_tbl1

n	\$epi
"s1":{"samp_id":1,"samp_data": { "episodes": [{"episodeID":10,"lengthMin":40,"minWat ched":25}, { "episodeID":20,"lengthMin":35,"minWatc hed":30}]}	{"episodeID":10,"lengthMin":40,"minWatc hed":25},

Table 6-6 (Cont.) Records in temp_tbl1

n	\$epi
"s1":{"samp_id":1,"samp_data":{"episodes":[{"episodeID":10,"lengthMin":40,"minWatched":25}, {"episodeID":20,"lengthMin":35,"minWatched":30}]}}	{"episodeID":20,"lengthMin":35,"minWatched":30}
"s1":{"samp_id":2,"samp_data":{"episodes":[{"episodeID":30,"lengthMin":40,"minWatched":25}, {"episodeID":40,"lengthMin":35,"minWatched":30}]}}	{"episodeID":30,"lengthMin":40,"minWatched":25},
"s1":{"samp_id":2,"samp_data":{"episodes":[{"episodeID":30,"lengthMin":40,"minWatched":25}, {"episodeID":40,"lengthMin":35,"minWatched":30}]}}	{"episodeID":40,"lengthMin":35,"minWatched":30}
"s1":{"samp_id":3,"samp_data":{"episodes":[{"episodeID":10,"lengthMin":40,"minWatched":25}, {"episodeID":20,"lengthMin":35,"minWatched":30}]}}	{"episodeID":10,"lengthMin":40,"minWatched":25},
"s1":{"samp_id":3,"samp_data":{"episodes":[{"episodeID":10,"lengthMin":40,"minWatched":25}, {"episodeID":20,"lengthMin":35,"minWatched":30}]}}	{"episodeID":20,"lengthMin":35,"minWatched":30}

Step 2 : Records in the second table temp_tbl2 - result of the full FROM clause (Above table joined with the result of the second path expression in the UNNEST clause (\$epi.values() as \$epiVal)

Table 6-7 Records in temp_tbl2

n	\$epi	\$epival
{"s1":{"samp_id":1,"samp_data":{"episodes":[{"episodeID":10,"lengthMin":40,"minWatched":25}, {"episodeID":20,"lengthMin":35,"minWatched":30}]}}	{"episodeID":10,"lengthMin":40,"minWatched":25},	10

Table 6-7 (Cont.) Records in temp_tbl2

n	\$epi	\$epival
{ "s1": { "samp_id":1, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":10, "lengthMin 40 ":40, "minWatched":25}	
{ "s1": { "samp_id":1, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":10, "lengthMin 25 ":40, "minWatched":25}	
{ "s1": { "samp_id":1, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":20, "lengthMin 20 ":35, "minWatched":30}	
{ "s1": { "samp_id":1, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":20, "lengthMin 35 ":35, "minWatched":30}	
{ "s1": { "samp_id":1, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":20, "lengthMin 30 ":35, "minWatched":30}	
{ "s1": { "samp_id":2, "samp_data": { "episodes": [{"episodeID":30, "lengthMi n":40, "minWatched":25}, { "episodeID":40, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":30, "lengthMin 30 ":40, "minWatched":25}	
{ "s1": { "samp_id":2, "samp_data": { "episodes": [{"episodeID":30, "lengthMi n":40, "minWatched":25}, { "episodeID":40, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":30, "lengthMin 40 ":40, "minWatched":25}	

Table 6-7 (Cont.) Records in temp_tbl2

n	\$epi	\$epival
{ "s1": { "samp_id":2, "samp_data": { "episodes": [{"episodeID":30, "lengthMi n":40, "minWatched":25}, { "episodeID":40, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":30, "lengthMin 25 ":40, "minWatched":25}	
{ "s1": { "samp_id":2, "samp_data": { "episodes": [{"episodeID":30, "lengthMi n":40, "minWatched":25}, { "episodeID":40, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":40, "lengthMin 40 ":35, "minWatched":30}	
{ "s1": { "samp_id":2, "samp_data": { "episodes": [{"episodeID":30, "lengthMi n":40, "minWatched":25}, { "episodeID":40, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":40, "lengthMin 35 ":35, "minWatched":30}	
{ "s1": { "samp_id":2, "samp_data": { "episodes": [{"episodeID":30, "lengthMi n":40, "minWatched":25}, { "episodeID":40, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":40, "lengthMin 30 ":35, "minWatched":30}	
{ "s1": { "samp_id":3, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":10, "lengthMin 10 ":40, "minWatched":25}	
{ "s1": { "samp_id":3, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":10, "lengthMin 40 ":40, "minWatched":25}	
{ "s1": { "samp_id":3, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":10, "lengthMin 25 ":40, "minWatched":25}	

Table 6-7 (Cont.) Records in temp_tbl2

n	\$epi	\$epival
{ "s1": { "samp_id":3, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":20, "lengthMin 20 ":35, "minWatched":30}	
{ "s1": { "samp_id":3, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":20, "lengthMin 35 ":35, "minWatched":30}	
{ "s1": { "samp_id":3, "samp_data": { "episodes": [{"episodeID":10, "lengthMi n":40, "minWatched":25}, { "episodeID":20, "lengthMin ":35, "minWatched":30}]]}	{ "episodeID":20, "lengthMin 30 ":35, "minWatched":30}	

Step 3 : There is no filter in the query and the fields in the SELECT clause are fetched. The final result of the query is :

Table 6-8 Query result

customer	episode_id	length	episode_details
2	30	25	30
2	30	25	40
2	30	25	25
2	40	30	40
2	40	30	35
2	40	30	30
1	10	25	10
1	10	25	40
1	10	25	25
1	20	30	20
1	20	30	35
1	20	30	30
3	10	25	10
3	10	25	40
3	10	25	25
3	20	30	20
3	20	30	35
3	20	30	30

Limitation for expression usage in the UNNEST clause

There are some restrictions on the expressions used in the UNNEST clause.

- All expressions used in the UNNEST clause must be path expressions.
- No predicates are allowed in any array-filter or map-filter steps within the path expressions.
- The expression must start with a variable. If the expression is a first expression in a list of expressions, then it should be a table alias associated with the target table.
- When you unnest nested arrays, each path expression unnests one level deeper, and it operates on the values produced by the previous level of unnesting. These values are represented by the variable associated with the previous path expression. So the starting variable of each expression must be the variable associated with the previous expression. This is not applicable if this is a first expression in a list of expressions in the UNNEST clause.
- The expression must finish with a `[]` or `.values()` step.
- The variables defined inside the UNNEST clause cannot be referenced in the rest of the FROM clause. They can be referenced within the UNNEST clause and outside the FROM clause.

For more details on path expression and to understand how these expressions are evaluated for an array or a map, see [Path Expressions](#).

Example: Using unnesting with a GROUP BY clause

Consider a TV streaming application. It streams various shows that are watched by customers across the globe. Every show has number of seasons and every season has multiple episodes. You need a persistent meta-data store which keeps track of the current activity of the customers using the TV streaming application. A customer is interested to know about the episodes they watched, the watch time per episode, the total number of seasons of the show they watched etc. The customer also wants the streaming application to start streaming from where they left off watching. The streaming application needs reports on which show is most popular among customers, how many minutes a show is being watched etc. These reports can be generated using UNNEST clause in queries.

To follow along with the examples, create the `stream_acct` table and insert data as described in the [Tables Used in the Examples](#) section.

Example 6-101 Fetch the different shows watched by people in the US alone and the number of people watching them

```
SELECT $show.showId, count(*) as cnt FROM stream_acct $s,  
unnest($s.acct_data.contentStreamed[] as $show)  
WHERE $s.acct_data.country = "USA" GROUP BY $show.showId  
ORDER BY count(*) DESC
```

Output:

```
{"showId":15,"cnt":2}  
{"showId":16,"cnt":2}
```

Example 6-102 For every show aired by the application, fetch the total watch time by all users

```
SELECT $show.showId, sum($show.seriesInfo.episodes.minWatched) AS total_time
FROM stream_acct $s, unnest($s.acct_data.contentStreamed[]) AS $show
GROUP BY $show.showId ORDER BY sum($show.seriesInfo.episodes.minWatched)
```

Output:

```
{"showId":26,"total_time":225}
{"showId":16,"total_time":440}
{"showId":15,"total_time":642}
```

Note

The unnest operator (that is the keyword unnest) can be omitted as it is a no-op operator. The use of the UNNEST clause is recommended when there is an index on the array(s) or map(s) that are being unnested. See [Examples: Using Indexes for Query Optimization](#) for more information.

The following query without the UNNEST clause is equivalent to the query above (with the UNNEST clause) and gives the same result.

```
SELECT $show.showId, sum($show.seriesInfo.episodes.minWatched) AS total_time
FROM stream_acct $s, $s.acct_data.contentStreamed[] AS $show
GROUP BY $show.showId ORDER BY sum($show.seriesInfo.episodes.minWatched)
```

Output:

```
{"showId":26,"total_time":225}
{"showId":16,"total_time":440}
{"showId":15,"total_time":642}
```

Example 6-103 Fetch the total watch time of users per show and season

```
SELECT $show.showId, $seriesInfo.seasonNum,
sum($seriesInfo.episodes.minWatched) AS length
FROM stream_acct n,
unnest(n.acct_data.contentStreamed[]) AS $show,
$show.seriesInfo[] as $seriesInfo)
GROUP BY $show.showId, $seriesInfo.seasonNum
ORDER BY sum($seriesInfo.episodes.minWatched)
```

Output:

```
{"showId":26,"seasonNum":2,"length":80}
{"showId":26,"seasonNum":1,"length":145}
{"showId":16,"seasonNum":2,"length":190}
{"showId":16,"seasonNum":1,"length":250}
```

```
{"showId":15,"seasonNum":2,"length":295}  
{"showId":15,"seasonNum":1,"length":347}
```

Example 6-104 Using unnest with a non-path expression

This is an example where the unnesting expression is not a path expression, and as a result, the UNNEST clause cannot be used. For example, a user may have multiple phone numbers in the same area code. To determine the number of users having phone numbers in different area codes, a particular user should be counted only once (even if the user has more than one phone number with the same area code). The below query returns, for each area code, the number of users having phone numbers in that area code. The query uses the `seq_distinct()` function, which eliminates duplicate values.

```
SELECT $area, count(*) AS cnt  
FROM Users u, seq_distinct(u.address.phones.area) AS $area  
GROUP BY $area
```

Output:

```
{"area":408,"cnt":1}  
{"area":831,"cnt":1}
```

To know more about `seq_distinct` function, see [Functions on Sequences](#).

Joins

Oracle NoSQL Database does not currently support the general join operators found in more traditional relational database systems. However, it does support a special kind of join among tables that belong to the same table hierarchy. These joins can be executed efficiently, because only co-located rows may match with each other. As a result, transferring very large amounts of data among servers is avoided.

A JOIN clause is used to combine rows from two or more tables, based on a related column between them. Joins are predominantly used in an Oracle NoSQL Database when a user is trying to extract data from tables that have a hierarchical relationship between them. For more details on hierarchical tables, see [Table Hierarchies](#).

There are three ways by which you can join tables in the same hierarchy in an Oracle NoSQL Database.

- NESTED TABLES
- Left Outer Join
- Inner Join

Using NESTED TABLES clause to query multiple tables in the same hierarchy

To query multiple tables in the same hierarchy, the NESTED TABLES clause must be used inside the FROM clause.

Syntax

```
nested_tables ::=
    NESTED TABLE "(" single_from_table
    [ANCESTORS "(" ancestor_tables ")"]
    [DESCENDANTS "(" descendant_tables ")"]
    ")"

ancestor_tables ::= nested_from_table ("," nested_from_table)*

descendant_tables ::= nested_from_table ("," nested_from_table)*

nested_from_table ::= aliased\_table\_name [ON or\_expression]
```

The NESTED TABLES clause specifies the participating tables and separates them in 3 groups. First the target table is specified. Then the ANCESTORS clause, if present, specifies a number of tables that must be ancestors of the target table in the table hierarchy. Finally, the DESCENDANTS clause, if present, specifies a number of tables that must be descendants of the target table in the table hierarchy. For each table an alias may be specified (and if not, one is created internally as described in the [CREATE TABLE Statement](#) section). The aliases must be unique.

Semantically, a NESTED TABLES clause is equivalent to a number of left-outer-join operations "centered" around the target table. The left-outer-join is an operation defined by standard SQL and supported by all major RDBMSs. For those not familiar with it already, we give a brief description in the last section of this chapter.

Our implementation of left outer join diverges slightly from the standard definition. The difference is in the "shape" of the results. Specifically, the result of a NESTED TABLES clause is a set of records, all having the same type, where (a) the number of fields is equal to the number of participating tables, (b) each field corresponds to one of the tables and stores either a row from that table or NULL, (c) the name of each field is the alias used by the associated table, and (d) the fields are ordered in the order that the participating tables would be encountered in a depth-first traversal of the table hierarchy.

So, in a NESTED TABLES result the columns of each table are grouped inside a subrecord. In contrast, the standard left-outer-join produces a "flat" result, where each result is a record/tuple whose number of fields is the sum of all the columns in the participating tables.

The mapping of a NESTED TABLES to a number of left-outer-joins is best explained with a few examples. For brevity, we will use the keyword LOJ in place of LEFT OUTER JOIN. Let's start with the following create table statements:

```
create table A (ida integer, a1 string, primary key(ida))

create table A.B (idb integer, b1 string, primary key(idb))

create table A.B.C (idc integer, c1 integer, primary key(idc))

create table A.B.C.D (idd integer, d1 double, primary key(idd))

create table A.B.E (ide integer, e1 integer, primary key(ide))

create table A.G (idg integer, g1 string, primary key(idg))

create table A.G.J (idj integer, j1 integer, primary key(idj))

create table A.G.H (idh integer, h1 integer, primary key(idh))

create table A.G.J.K (idk integer, k1 integer, primary key(idk))
```

The above statements create the following table hierarchy:

The NESTED TABLES clause specifies the join tree as a "projection" of the table hierarchy that includes the tables in the NESTED TABLES. For example, the join tree for NESTED TABLES(A.B) ancestors(A) descendants (A.B.C.D, A.B.E) is shown below. The arrows indicate the direction of the LOJs (from the left table to the right table).

Now, let's look at the following NESTED TABLES cases and their equivalent LOJ operations

1. NESTED TABLES (A.B.C c ancestors(A a, A.B b))

is equivalent to

```
A.B.C c LOJ A a ON c.ida = a.ida LOJ A.B b ON c.ida = b.ida AND
c.idb = b.idb
```

We can see that the join predicates are implicit in the NESTED TABLES clause, and they are always on the primary key columns of the participating tables.

Because for each A.B.C row there is at most one matching A and A.B row, the number of records in the result is the same as the number of A.B.C rows. This is always true when the NESTED TABLES clause includes ancestors only. In this case, the effect of the operation is to *decorate* the target table rows with the columns from the matching ancestor rows (if any), without eliminating or adding any other rows.

2. NESTED TABLES (A a descendants(A.B b, A.B.C c))

is equivalent to

```
A a LOJ A.B b ON a.ida = b.ida
  LOJ A.B.C c ON b.ida = c.ida AND b.idb = c.idb
```

Another way to explain the semantics of the DESCENDANTS clause is to use the contains relationship defined in [Table Hierarchies](#) section, but restricted to the descendant tables in the join tree only. Let R be a target table row, and S(R) be the set containing all the descendant rows that are reachable from R via the contains relationship (i.e., S(R) is the transitive closure of contains applied on R). If S(R) is empty, a single record is returned for R, that contains R and a NULL value for each of the descendant tables in the join tree. Otherwise, let B(R) be the boundary subset of S(R), i.e., all rows in S(R) that do not have any descendant rows in the join tree. Then, a result is generated for each row in B(R) as follows: Let RR be a row in B(R) and T be its containing table. The result associated with RR is a record containing all the rows in the path from R to RR and a NULL value for every table that is not in the path from the target table to T.

3. NESTED TABLES (A a descendants(A.B b, A.G g))

is equivalent to

```
A a LOJ
  (A.B b UNION A.G g)
  ON (a.ida = b.ida or b.ida IS NULL) and (a.ida = g.ida or g.ida IS NULL)
```

As in case 2, target table A is joined with 2 descendant tables. However, because the descendant tables come from 2 different branches of the join tree, we have to use a UNION operation in the SQL expression above. This UNION unions both the rows and the columns of tables A.B and A.G. So, if table A.B has N rows with n columns each and table A.G has M rows with m columns, the result of the UNION has N + M rows, with n + m columns each. The first N rows contain the rows of A.B with NULL values for the m columns, and the following M rows contain the rows of A.G with NULL values for the n columns. When matching A rows with the UNION rows, we distinguish whether a UNION row comes from table A.B or A.g by checking whether g.ida is NULL or b.ida is NULL, respectively.

Notice that the contains-base definition given in case 2 applies here as well.

4. NESTED TABLES (A a descendants(A.B b, A.B.C c, A.B.E e, A.G.J.K k))

is equivalent to

```
A a LOJ
  (
    A.B b LOJ
      (A.B.C c UNION A.B.E e)
      ON (b.ida = c.ida and b.idb = c.idb or c.ida IS NULL) and
        (b.ida = e.ida and b.idb = e.idb or e.ida IS NULL)
    UNION
    A.G.J.K k
```

```
)
ON (a.ida = b.ida or b.ida IS NULL) and (a.ida = k.ida or k.ida IS NULL)
```

This example is just a more complex version of case 3.

5. NESTED TABLES (A.B b ancestors(A a ON a.a1 = "abc")
descendants(A.B.C c ON c.c1 > 10,
A.B.C.D d,
A.B.E e))

is equivalent to

```
(A.B b LOJ A a ON b.ida = a.ida and a.a1 = "abc") LOJ
(
  A.B.C c LOJ A.B.C.D
  ON c.ida = d.ida and c.idb = d.idb and c.idc = d.idc
  UNION
  E
)
ON (b.ida = c.ida and b.idb = c.idb or c.ida IS NULL) and
(b.ida = e.ida and b.idb = e.idb or e.ida IS NULL)
```

This is an example that just puts everything together. It contains both ANCESTOR and DESCENDANT clauses, as well as ON predicates. The mapping from NESTED TABLES to LOJs uses the same patterns as in the previous cases. The ON predicates are just added with the join predicates. In most cases, the ON predicates inside a NESTED TABLES will be local predicates on the right table of an LOJ, but more generally, they can also reference any columns from any table that is an ancestor of the table the ON appears next to.

Example: Using NESTED TABLES clause to query multiple tables in the same hierarchy

Let's consider an application that tracks a population of users and the emails sent or received by these users. Given that SQL for Oracle NoSQL Database does not currently support general purpose joins, the emails are stored in a table that is created as a child of users, so that queries can be written that combine information from both tables using the NESTED TABLES clause. The create table statements for the two tables are shown below. Notice that it is possible for the application to receive emails that are not associated with any user in the users table; such emails will be assigned a uid value that does not exist in the users table.

```
CREATE TABLE users(
  uid INTEGER,
  name string,
  email_address string,
  salary INTEGER,
  address json,
  PRIMARY KEY(uid))

CREATE TABLE users.emails(
  eid INTEGER,
  sender_address string, // sender email address
```

```
receiver_address string, // receiver email address
time timestamp(3),
size INTEGER,
content string,
PRIMARY KEY(eid))
```

Here are two queries that can be written over the users and emails tables.

Example 6-105 Joining Tables

Count the number of emails sent in 2017 by all users whose salary is greater than 200K

```
SELECT count(eid)
FROM NESTED TABLES(
  users
  descendants(users.emails ON email_address = sender_address and
    year(time) = 2017)
)
WHERE salary > 200
```

In the above query, we are using count(eid) rather than count(*) because there may exist users with no emails or no sent emails. For such a user, the FROM clause will return a record where the eid field will be NULL, and count(eid) will not count such a record.

Example 6-106 Joining Tables

For each email whose size is greater than 100KB and was sent by a user in the the users table, return the name and address of that user.

```
SELECT name, address
FROM NESTED TABLES(users.emails ancestors(users))
WHERE size > 100 AND sender_address = email_address
```

In the above query will return duplicate results for any user who has sent more than one "large" email. Currently, SQL for Oracle NoSQL Database does not support SELECT DISTINCT, so the duplicate elimination has to be performed by the application.

Left Outer Join (LOJ)

A Left Outer Join (LOJ) is one of the join operations that allows you to specify a join clause. It preserves the unmatched rows from the first (left) table, joining them with a NULL row in the second (right) table. This means all left rows that do not have a matching row in the right table will appear in the result, paired with a NULL value in place of a right row.

Syntax of LOJ

```
loj_from_clause ::= FROM ( aliased\_table\_name | left\_outer\_join\_tables
left_outer_join_table ::= LEFT OUTER JOIN single_from_table ON expression
left_outer_join_tables ::= single_from_table left_outer_join_table
(left_outer_join_table)*
```

Semantics

The FROM clause specifies the participating tables and separates them into two groups. First, the target table (the table which is on the left side of the LEFT OUTER JOIN clause) is

specified. Then a LEFT OUTER JOIN clause can be specified. The table to the left of the LEFT OUTER JOIN keywords is called the left table, and the one to the right of LEFT OUTER JOIN is the right table. The name of any participating table in the join may be followed by a table alias. A table alias is a temporary name given to a table. Aliases are often used to make column names more readable. If no alias is specified, one is created internally, using the name of the table as it is spelled in the query, but with the dot(".") character replaced with '_' in the case of child tables. A join predicate specifies the columns on which records from two or more tables are joined. The expression after the ON clause lists all the join predicates between the two tables. The `left_outer_join_tables` clause specifies that the result of one LOJ can be the target table for another LOJ.

Like other kinds of joins, the LOJ creates a result set containing pairs of matching rows from the left and right tables. However, an LOJ will also preserve all rows of the left table, that is, a left row that does not have a matching row will appear in the result, paired with a NULL value in place of a right row.

Result set ordering in an LOJ:

In an LOJ, the order of fields in the result-set is always in top-down order. That means the order of output in the result set is always from the ancestor table first and then the descendant table. This is true irrespective of the order of the joins.

Different scenarios of using an LOJ

To understand different scenarios of using an LOJ, let us create a few hierarchical tables and populate them with some data.

```
CREATE TABLE IF NOT EXISTS A (  
    ida INTEGER,  
    a1 string,  
    PRIMARY KEY(ida))
```

```
CREATE TABLE IF NOT EXISTS A.B (  
    idb INTEGER,  
    b1 string,  
    PRIMARY KEY(idb))
```

```
CREATE TABLE IF NOT EXISTS A.B.C (  
    idc INTEGER,  
    c1 string,  
    PRIMARY KEY(idc))
```

```
CREATE TABLE IF NOT EXISTS A.G (  
    idg INTEGER,
```

```
g1 string,  
PRIMARY KEY(idg))
```

```
INSERT INTO A VALUES(1, 'a1')
```

```
INSERT INTO A VALUES(2, 'a2')
```

```
INSERT INTO A VALUES(3, 'a3')
```

```
INSERT INTO A.B VALUES(1, 1, 'b1')
```

```
INSERT INTO A.B VALUES(2, 2, 'b2')
```

```
INSERT INTO A.B.C VALUES(1, 1, 1, 'c1')
```

```
INSERT INTO A.B.C VALUES(2, 2, 1, 'c2')
```

```
INSERT INTO A.B.C VALUES(3, 3, 1, 'c3')
```

```
INSERT INTO A.G VALUES(1, 1, 'g1')
```

```
INSERT INTO A.G VALUES(2, 2, 'g2')
```

Join with Descendants

You can join a target table with its descendant. All rows of the target table(left table) will be returned and for those rows where there is no match in the descendant's table(right table), NULL values are populated as shown in the examples below.

Example 6-107 A target table A is joined with its child table A.B

```
SELECT * FROM A a LEFT OUTER JOIN A.B b  
ON a.ida = b.ida ORDER BY a.ida
```

Output

```
{"a":{"ida":1,"a1":"a1"},"b":{"ida":1,"idb":1,"b1":"b1"}}  
{"a":{"ida":2,"a1":"a2"},"b":{"ida":2,"idb":2,"b1":"b2"}}  
{"a":{"ida":3,"a1":"a3"},"b":null}
```

Example 6-108 A target table A is joined with its descendant A.B.C

```
SELECT * FROM A a LEFT OUTER JOIN A.B.C c  
ON a.ida = c.ida ORDER BY a.ida
```

Output

```
{ "a": { "ida": 1, "a1": "a1" }, "c": { "ida": 1, "idb": 1, "idc": 1, "c1": "c1" } }
{ "a": { "ida": 2, "a1": "a2" }, "c": { "ida": 2, "idb": 2, "idc": 1, "c1": "c2" } }
{ "a": { "ida": 3, "a1": "a3" }, "c": { "ida": 3, "idb": 3, "idc": 1, "c1": "c3" } }
```

Join with Ancestors

You can join a target table with its ancestor. All rows of the target table(left table) will be returned and for those rows where there is no match in the ancestor table(right table), NULL values are populated as shown in the examples below.

Example 6-109 A target table A.B is joined with its parent table A

```
SELECT * FROM A.B b LEFT OUTER JOIN A a
ON a.ida = b.ida ORDER BY a.ida
```

Output

```
{ "a": { "ida": 1, "a1": "a1" }, "b": { "ida": 1, "idb": 1, "b1": "b1" } }
{ "a": { "ida": 2, "a1": "a2" }, "b": { "ida": 2, "idb": 2, "b1": "b2" } }
```

Example 6-110 A target table A.B.C is joined with its ancestor A.B

```
SELECT * FROM A.B.C c LEFT OUTER JOIN A.B b
ON c.ida = b.ida and c.idb = b.idb
```

Output

```
{ "b": null, "c": { "ida": 3, "idb": 3, "idc": 1, "c1": "c3" } }
{ "b": { "ida": 2, "idb": 2, "b1": "b2" }, "c": { "ida": 2, "idb": 2, "idc": 1, "c1": "c2" } }
{ "b": { "ida": 1, "idb": 1, "b1": "b1" }, "c": { "ida": 1, "idb": 1, "idc": 1, "c1": "c1" } }
```

Multiple LOJs in one SELECT statement

A target table (A) can be joined to its child (A.B) and the result of the join can be joined to the descendant (A.B.C) of the target table. If the right table is a descendant table, the ON expression should contain join predicates on all the primary key columns of the left table using an AND clause. In the example given below, the left table has two primary key columns (ida and idb). An AND clause is used with these two join predicates (b.ida = c.ida AND b.idb = c.idb).

Example 6-111 Multiple LOJs of the parent table with its descendants

```
SELECT * FROM A a LEFT OUTER JOIN A.B b ON a.ida = b.ida
LEFT OUTER JOIN A.B.C c ON b.ida = c.ida AND b.idb = c.idb ORDER BY a.ida
```

Output

```
{ "a": { "ida": 1, "a1": "a1" }, "b": { "ida": 1, "idb": 1, "b1": "b1" }, "c":
{ "ida": 1, "idb": 1, "idc": 1, "c1": "c1" } }
{ "a": { "ida": 2, "a1": "a2" }, "b": { "ida": 2, "idb": 2, "b1": "b2" }, "c":
```

```
{ "ida":2, "idb":2, "idc":1, "c1":"c2" } }
{ "a":{ "ida":3, "a1":"a3" }, "b":null, "c":null }
```

Example 6-112 Multiple LOJs of the parent table with its ancestors

```
SELECT * FROM A.B.C c LEFT OUTER JOIN A a ON c.ida = a.ida
LEFT OUTER JOIN A.B b ON c.ida = b.ida AND c.idb = b.idb ORDER BY c.ida, c.idb
```

Output

```
{ "a":{ "ida":1, "a1":"a1" }, "b":{ "ida":1, "idb":1, "b1":"b1" }, "c":
{ "ida":1, "idb":1, "idc":1, "c1":"c1" } }
{ "a":{ "ida":2, "a1":"a2" }, "b":{ "ida":2, "idb":2, "b1":"b2" }, "c":
{ "ida":2, "idb":2, "idc":1, "c1":"c2" } }
{ "a":{ "ida":3, "a1":"a3" }, "b":null, "c":{ "ida":3, "idb":3, "idc":1, "c1":"c3" } }
```

Join with an ancestor and a descendant

You can also join a target table to an ancestor and a descendant. First, join the target table to its ancestor and then join the result of this query with the descendant of the target table. Any number of such joins is possible in a single SQL statement. In all the joins, all rows of the left table will be returned and for those rows where there is no match in the right table, NULL values are populated. If the right table is a descendant table, the ON expression should contain join predicates on all the primary key columns of the left table using an "AND" clause.

Example 6-113 A target table A.B is joined with its ancestor A, the result of which is joined with its descendant A.B.C

```
SELECT * FROM A.B b LEFT OUTER JOIN A a ON b.ida = a.ida AND a.a1 = "abc"
LEFT OUTER JOIN A.B.C c ON b.ida = c.ida AND b.idb = c.idb
```

Output

```
{ "a":null, "b":{ "ida":1, "idb":1, "b1":"b1" }, "c":
{ "ida":1, "idb":1, "idc":1, "c1":"c1" } }
{ "a":null, "b":{ "ida":2, "idb":2, "b1":"b2" }, "c":
{ "ida":2, "idb":2, "idc":1, "c1":"c2" } }
```

Non-join predicate restrictions in an LOJ

An LOJ can additionally have non-join predicates (that is the columns that are not part of the join) as a restriction to filter data. The result set of an LOJ depends on whether you place the non-join predicate restriction in the ON clause or in the WHERE clause. The non-join predicate in an ON clause is just applied to the join operation, whereas the non-join predicate in the WHERE clause will apply to the entire results-set.

Example 6-114 Non-join predicate in the WHERE clause

If a non-join predicate is placed in the WHERE clause, the restriction is applied to the result of the outer join. That is, it removes all the rows for which the WHERE condition is not TRUE. In the example below, you get only one row that matches the WHERE condition as the result set.

```
SELECT * FROM A a LEFT OUTER JOIN A.B b ON a.ida = b.ida
WHERE b.ida > 1 ORDER BY a.ida
```

Output

```
{"a":{"ida":2,"a1":"a2"},"b":{"ida":2,"idb":2,"b1":"b2"}}
```

Example 6-115 Non-join predicate in the ON clause

If you move the non-join predicate restriction to the ON clause, the result set includes all the rows that meet the ON clause condition. Rows from the right outer table that do not meet the ON condition are populated with NULL values as shown below.

```
SELECT * FROM A a LEFT OUTER JOIN A.B b ON a.ida = b.ida
AND b.ida > 1 ORDER BY a.ida
```

Output

```
{"a":{"ida":1,"a1":"a1"},"b":null}
{"a":{"ida":2,"a1":"a2"},"b":{"ida":2,"idb":2,"b1":"b2"}}
{"a":{"ida":3,"a1":"a3"},"b":null}
```

Limitations of LOJ

There are some limitations while using an LOJ in an Oracle NoSQL Database.

Joins with Siblings

LOJs in an Oracle NoSQL Database cannot be applied to siblings. LOJ can be used in tables that only have a direct hierarchical relationship.

Example 6-116 Error scenario: LOJ on siblings

```
SELECT * FROM A.B b LEFT OUTER JOIN A.G g ON b.ida=g.ida
```

Output

```
Error handling command select * from A.B b LEFT OUTER JOIN A.G g on
b.ida=g.ida:
Error: at (1, 40) Table A.G is neither ancestor nor descendant of the target
table A.B
```

Order of tables in an LOJ

While using multiple joins in a single statement, the tables in an LOJ must appear in top-down order after the target table. A child table cannot be to the left of its parent in a multiple join operation.

Example 6-117 Multiple LOJ of the parent table with its ancestors in the wrong order

In this example, an error is thrown as the parent table A cannot be at the right of its descendant A.B.

```
SELECT * FROM A.B.C c LEFT OUTER JOIN A.B b ON c.ida = b.ida and c.idb =
b.idb
LEFT OUTER JOIN A a on c.ida = a.ida
```

Output

Error handling command `SELECT * FROM A.B.C c LEFT OUTER JOIN A.B b ON c.ida = b.ida`
and `c.idb = b.idb LEFT OUTER JOIN A a ON c.ida = a.ida`:
Error: at (3, 22) Table A is not descendant of table A.B.
Tables in left-outer-joins must appear in top-down order after the target table

Join predicates

In an LOJ, the ON expression of the right table must contain all the required join predicates that link the parent and child tables. Missing any join predicate results in an exception.

Example 6-118 If the right table is an ancestor of the target table, the ON expression must contain join-predicates on all the primary key columns of the right table

In this example, one of the join predicates on the primary key column `idb` of the right table is missing which throws an error.

```
SELECT * FROM A.B.C c LEFT OUTER JOIN A a ON c.ida = a.ida
LEFT OUTER JOIN A.B b ON c.ida = b.ida
```

Output

Error handling command `SELECT * FROM A.B.C c LEFT OUTER JOIN A a ON c.ida = a.ida`
`LEFT OUTER JOIN A.B b ON c.ida = b.ida`:
Error: A join predicate is missing from ON clause of table A.B : `c.idb = b.idb`

Example 6-119 If the right table is a descendant of the target table, the ON expression must contain join-predicates on all the primary key columns of the left table

In this example, one of the join predicates on the primary key column `idb` of the left table is missing which throws an error.

```
SELECT * FROM A.B b LEFT OUTER JOIN A.B.C c ON b.ida = c.ida
```

Output

Error handling command `SELECT * FROM A.B b LEFT OUTER JOIN A.B.C c ON b.ida = c.ida`:
Error: A join predicate is missing from ON clause of table A.B.C : `b.idb = c.idb`

Nested tables Vs LOJ

To query multiple tables in the same hierarchy, you can also use the `NESTED TABLES` clause. To get more details about Nested tables and using it to join a parent table with its child table, see [Using NESTED TABLES clause to query multiple tables in the same hierarchy](#)

Table 6-9 Nested Tables Vs LOJ

Nested Tables	LOJ
Queries multiple tables in the same hierarchy	Queries multiple tables in the same hierarchy
Not an ANSI-SQL Standard	ANSI-SQL Standard
Supports sibling tables join	Does not support sibling table joins

Example: Using Left Outer Joins

Let us take an example of an online shopping portal that uses the Oracle NoSQL Database. This application has millions of customers who do online shopping and add things they want to buy in a shopping cart. Finally, the order is processed and billed. The shopping pattern of customers can be analyzed using this application.

Here there are three tables in a parent-child relationship. The customerprofile table with all the personal information of customers is the top-level table. Shoppingcart which lists the items picked by the customer, their quantity, and the individual price of each item is the descendant table of customerprofile. Order which has the total order quantity and the total order value. is the descendant table of shoppingcart.

Let us create these tables and populate some data into them.

```
CREATE TABLE customerprofile (
  customer_id INTEGER,
  customer_name STRING,
  customer_address STRING,
  customer_phone INTEGER,
  PRIMARY KEY(customer_id))
```

```
CREATE TABLE customerprofile.shoppingcart (
  cart_id STRING,
  product_name STRING,
  product_quantity INTEGER,
  product_price INTEGER,
  PRIMARY KEY(cart_id))
```

```
CREATE TABLE customerprofile.shoppingcart.order (
  order_id INTEGER,
  order_quantity INTEGER,
  order_total INTEGER,
  PRIMARY KEY(order_id))
```

```
INSERT INTO customerprofile VALUES (1,"Aana","Blr",111111111)
INSERT INTO customerprofile VALUES (2,"Bobby","Chn",22222222)
INSERT INTO customerprofile VALUES (3,"Celin","Del",33333333)
INSERT INTO customerprofile VALUES (4,"Diana","Blr",44444444)
INSERT INTO customerprofile VALUES (5,"Elizabeth","Mum",55555555)
```

```
INSERT INTO customerprofile.shoppingcart VALUES(1,'c1',"Shampoo",2,300)
INSERT INTO customerprofile.shoppingcart VALUES(1,'c2',"Soap",3,80)
```

```
INSERT INTO customerprofile.shoppingcart VALUES(2, 'c3', "Milk", 5, 100)
INSERT INTO customerprofile.shoppingcart VALUES(3, 'c4', "Chips", 2, 50)
INSERT INTO customerprofile.shoppingcart VALUES(4, 'c5', "Bread", 1, 40)
```

```
INSERT INTO customerprofile.shoppingcart.order VALUES(1, 'c1', 100, 2, 600)
INSERT INTO customerprofile.shoppingcart.order VALUES(1, 'c2', 101, 3, 240)
INSERT INTO customerprofile.shoppingcart.order VALUES(2, 'c3', 102, 5, 500)
INSERT INTO customerprofile.shoppingcart.order VALUES(3, 'c4', 103, 2, 100)
```

Examples using LOJ

Example 6-120 Use case 1: Fetch all customer details and their order history

```
SELECT * FROM customerprofile a LEFT OUTER JOIN
customerprofile.shoppingcart.order c
ON a.customer_id = c.customer_id ORDER BY a.customer_id
```

Output:

```
{"a":
{"customer_id":1,"customer_name":"Aana","customer_address":"Blr","customer_phone":111111111}, "c":
{"customer_id":1,"cart_id":"c1","order_id":100,"order_quantity":2,"order_total":600}}
{"a":
{"customer_id":1,"customer_name":"Aana","customer_address":"Blr","customer_phone":111111111}, "c":
{"customer_id":1,"cart_id":"c2","order_id":101,"order_quantity":3,"order_total":240}}
{"a":
{"customer_id":2,"customer_name":"Bobby","customer_address":"Chn","customer_phone":22222222}, "c":
{"customer_id":2,"cart_id":"c3","order_id":102,"order_quantity":5,"order_total":500}}
{"a":
{"customer_id":3,"customer_name":"Celin","customer_address":"Del","customer_phone":33333333}, "c":
{"customer_id":3,"cart_id":"c4","order_id":103,"order_quantity":2,"order_total":100}}
{"a":
{"customer_id":4,"customer_name":"Diana","customer_address":"Blr","customer_phone":44444444}, "c":null}
{"a":
{"customer_id":5,"customer_name":"Elizabeth","customer_address":"Mum","customer_phone":55555555}, "c":null}
```

Example 6-121 Use case 2: Fetch all customers who have shopped for at least 3 pieces of the same item

```
SELECT * FROM customerprofile a LEFT OUTER JOIN customerprofile.shoppingcart b
ON a.customer_id = b.customer_id WHERE product_quantity >2
```


Output:

```
{
  "a": {
    "customer_id": 1, "customer_name": "Aana", "customer_address": "Blr", "customer_phone": 1111111111, "b": {
      "customer_id": 1, "cart_id": "c2", "product_name": "Soap", "product_quantity": 3, "product_price": 80
    }
  },
  "a": {
    "customer_id": 2, "customer_name": "Bobby", "customer_address": "Chn", "customer_phone": 2222222222, "b": {
      "customer_id": 2, "cart_id": "c3", "product_name": "Milk", "product_quantity": 5, "product_price": 100
    }
  }
}
```

Example 6-122 Use case 3: Fetch the details of only those customers who have at least shopped for a value of 500

```
SELECT * FROM customerprofile a LEFT OUTER JOIN
customerprofile.shoppingcart.order b
ON a.customer_id = b.customer_id WHERE order_total >=500
```

Output:

```
{
  "a": {
    "customer_id": 1, "customer_name": "Aana", "customer_address": "Blr", "customer_phone": 1111111111,
    "b": {
      "customer_id": 1, "cart_id": "c1", "order_id": 100, "order_quantity": 2, "order_total": 600
    }
  },
  "a": {
    "customer_id": 2, "customer_name": "Bobby", "customer_address": "Chn", "customer_phone": 2222222222,
    "b": {
      "customer_id": 2, "cart_id": "c3", "order_id": 102, "order_quantity": 5, "order_total": 500
    }
  }
}
```

Inner Join

An inner join is an operation that produces new rows by combining rows from two or more tables, based on the join predicates applied to related columns or fields between them. The result-set contains only those combined rows that satisfy the join predicates.

Syntax

The tables to be joined are specified in the FROM clause of the SELECT statement, separated by commas.

```
innerjoin_from_clause ::= FROM innerjoin\_tables (COMMA innerjoin\_tables)* (COMMA unnest\_syntax)*
```

```
innerjoin_table_clause ::= single_from_table | aliased\_table\_name | left\_outer\_join\_tables | nested\_tables
```

For description of various elements of the syntax, see [FROM](#) clause. Note that the syntax implies that we can not only join individual tables, but also the result of a [NESTED TABLES](#) clause (or [left outer join](#) list) with an individual table, or with the result of another NESTED TABLES clause (or left outer join list).

Semantics

The FROM clause specifies the participating tables. The name of any participating table in the join may be followed by a table alias.

Let us consider that we perform an inner join of two tables A and B. Table A has N rows and n columns, and table B has M rows and m columns. Conceptually, every row in table A is joined with every row in table B. The resultant table AB would have $N * M$ rows and $n + m$ columns.

Let us understand how inner join works using a simple example. In the below example table A has 3 rows, 3 columns with CA1 as the primary key/shard key; and its child table B has 2 rows, 3 columns with CB1 as the primary key.

Table 6-10 Table A

CA1	CA2	CA3
RA ₁₁	RA ₁₂	RA ₁₃
RA ₂₁	RA ₂₂	RA ₂₃
RA ₃₁	RA ₃₂	RA ₃₃

Table 6-11 Table B

(CA1)	CB1	CB2	CB3
RA ₁₁	RB ₁₁	RB ₁₂	RB ₁₃
RA ₂₁	RB ₂₁	RB ₂₂	RB ₂₃

Column (CA1) is the inherited shard key from table A.

Table 6-12 Combined Table AB (every row in table A is combined with every row in table B)

CA1	CA2	CA3	(CA1)	CB1	CB2	CB3
RA₁₁	RA ₁₂	RA ₁₃	RA₁₁	RB ₁₁	RB ₁₂	RB ₁₃
RA ₁₁	RA ₁₂	RA ₁₃	RA ₂₁	RB ₂₁	RB ₂₂	RB ₂₃
RA ₂₁	RA ₂₂	RA ₂₃	RA ₁₁	RB ₁₁	RB ₁₂	RB ₁₃
RA₂₁	RA ₂₂	RA ₂₃	RA₂₁	RB ₂₁	RB ₂₂	RB ₂₃
RA ₃₁	RA ₃₂	RA ₃₃	RA ₁₁	RB ₁₁	RB ₁₂	RB ₁₃
RA ₃₁	RA ₃₂	RA ₃₃	RA ₂₁	RB ₂₁	RB ₂₂	RB ₂₃

The WHERE clause is then applied to this combined table. The WHERE clause contains the join predicates that specifies the columns or fields on which the records from the tables are joined.

In the case of inner join, the WHERE clause must include the equality predicate on all shard keys of the participating tables. Therefore, after applying this predicate, the final table after inner join would be:

Table 6-13 Result after inner join (based on shard key equality predicate)

CA1	CA2	CA3	(CA1)	CB1	CB2	CB3
RA ₁₁	RA ₁₂	RA ₁₃	RA ₁₁	RB ₁₁	RB ₁₂	RB ₁₃
RA ₂₁	RA ₂₂	RA ₂₃	RA ₂₁	RB ₂₁	RB ₂₂	RB ₂₃

If there are other join predicates in the WHERE clause, those are then applied.

When joining more than 2 tables, say A, B, C, and D, the FROM clause is applied first to tables A and B producing table AB. In a similar manner, the FROM clause is applied to tables AB and C, producing table ABC. Lastly, the FROM clause is applied to ABC and D, producing table ABCD. The join predicates in the WHERE clause are then applied to this combined table ABCD.

While performing an inner join, the following are applicable:

- Only joins among tables in the same [table hierarchy](#) are allowed.
- Supports joining of tables that are in an ancestor-descendant relationship as well as tables that are not in an ancestor-descendant relationship.
- The join predicates must include equality predicates between all the shard key columns of the joined tables. That is, for any pair of joined tables, a row from one table will match with a row from the other table only if they both have the same values on their shard key columns. To understand more about shard keys, see [CREATE TABLE](#). You can use the [DESCRIBE TABLE](#) statement to identify the shard keys.
- The rest of the predicates in the WHERE clause are applied to these matching rows.

Ordering of columns in the result-set of inner join:

If you use a '*' with the SELECT clause, wherein all the fields in the tables are returned, the order of fields in the result-set depends on the order in which you specify the tables in the FROM clause. If you provide a list of fields in the SELECT clause, then the order of the fields is as specified in the SELECT clause.

Different Scenarios of Using an Inner Join

To understand the different scenarios of using an inner join, let us create a few hierarchical tables and populate them with some data.

```
CREATE TABLE IF NOT EXISTS A (
    ida INTEGER,
    a1 string,
    PRIMARY KEY(ida))

CREATE TABLE IF NOT EXISTS A.B (
    idb INTEGER,
    b1 string,
    PRIMARY KEY(idb))

CREATE TABLE IF NOT EXISTS A.B.C (
    idc INTEGER,
    c1 string,
    PRIMARY KEY(idc))

CREATE TABLE IF NOT EXISTS A.G (
```

```

        idg INTEGER,
        g1 string,
        PRIMARY KEY(idg))

INSERT INTO A VALUES(1, 'a1')
INSERT INTO A VALUES(2, 'a2')
INSERT INTO A VALUES(3, 'a3')
INSERT INTO A.B VALUES(1, 1, 'b1')
INSERT INTO A.B VALUES(2, 2, 'b2')
INSERT INTO A.B.C VALUES(1, 1, 1, 'c1')
INSERT INTO A.B.C VALUES(2, 2, 1, 'c2')
INSERT INTO A.B.C VALUES(3, 3, 1, 'c3')
INSERT INTO A.G VALUES(1, 1, 'g1')
INSERT INTO A.G VALUES(2, 2, 'g2')

```

Join with Descendant Table:**Example 1a:**

Parent table A is joined with descendant table B. Note that only the matching rows are returned. The unmatched row from table A with ida 3, is not returned in the output.

```
SELECT * FROM A a, A.B b WHERE a.ida = b.ida ORDER BY a.ida
```

Output:

```

{"a":{"ida":1,"a1":"a1"},"b":{"ida":1,"idb":1,"b1":"b1"}}
{"a":{"ida":2,"a1":"a2"},"b":{"ida":2,"idb":2,"b1":"b2"}}

```

2 rows returned

Example 1b:

Parent table A is joined with descendant table C. A descendant table can be any level hierarchically below a table (For example C is the child of B which is the child of A, so C is a descendant of A).

```
SELECT * FROM A a, A.B.C c WHERE a.ida = c.ida ORDER BY a.ida
```

Output:

```

{"a":{"ida":1,"a1":"a1"},"c":{"ida":1,"idb":1,"idc":1,"c1":"c1"}}
{"a":{"ida":2,"a1":"a2"},"c":{"ida":2,"idb":2,"idc":1,"c1":"c2"}}
{"a":{"ida":3,"a1":"a3"},"c":{"ida":3,"idb":3,"idc":1,"c1":"c3"}}

```

3 rows returned

Join with Ancestor Table:**Example 1a:**

Table B is joined with ancestor table A.

```
SELECT * FROM A.B b, A a WHERE a.ida = b.ida ORDER BY a.ida
```

Output:

```
{ "b": { "ida": 1, "idb": 1, "b1": "b1" }, "a": { "ida": 1, "a1": "a1" } }
{ "b": { "ida": 2, "idb": 2, "b1": "b2" }, "a": { "ida": 2, "a1": "a2" } }
```

2 rows returned

Multiple Table Join:**Example 1:**

Table C is joined with ancestor tables A and B.

```
SELECT * FROM A.B.C c, A a ,A.B b WHERE c.ida = b.ida AND c.ida = a.ida ORDER
BY c.ida, c.idb
```

Output:

```
{ "c": { "ida": 1, "idb": 1, "idc": 1, "c1": "c1" }, "a": { "ida": 1, "a1": "a1" }, "b":
{ "ida": 1, "idb": 1, "b1": "b1" } }
{ "c": { "ida": 2, "idb": 2, "idc": 1, "c1": "c2" }, "a": { "ida": 2, "a1": "a2" }, "b":
{ "ida": 2, "idb": 2, "b1": "b2" } }
```

2 rows returned

Sibling Join:**Example 1:**

Table B is joined with sibling table G. Both the tables are children of table A.

```
SELECT * FROM A.B b,A.G g WHERE b.ida=g.ida
```

Output:

```
{ "b": { "ida": 2, "idb": 2, "b1": "b2" }, "g": { "ida": 2, "idg": 2, "g1": "g2" } }
{ "b": { "ida": 1, "idb": 1, "b1": "b1" }, "g": { "ida": 1, "idg": 1, "g1": "g1" } }
```

2 rows returned

Join with Ancestor and Descendant Tables:**Example 1:**

Table B is joined with ancestor A and descendant C.

```
SELECT * FROM A.B b ,A a ,A.B.C c WHERE c.ida = a.ida AND b.ida = a.ida
```

Output:

```
{ "b": { "ida": 1, "idb": 1, "b1": "b1" }, "a": { "ida": 1, "a1": "a1" }, "c":
{ "ida": 1, "idb": 1, "idc": 1, "c1": "c1" } }
{ "b": { "ida": 2, "idb": 2, "b1": "b2" }, "a": { "ida": 2, "a1": "a2" }, "c":
{ "ida": 2, "idb": 2, "idc": 1, "c1": "c2" } }
```

2 rows returned

Join with Table that is neither an Ancestor nor a Descendant:

Example 1:

Table G is joined with C. Parent of table G is A and parent of table C is B. Table G and C are not in an ancestor-descendant relationship.

```
SELECT * FROM A.G g, A.B.C c WHERE g.ida=c.ida
```

Output:

```
{"g":{"ida":1,"idg":1,"g1":"g1"},"c":{"ida":1,"idb":1,"idc":1,"c1":"c1"}}
{"g":{"ida":2,"idg":2,"g1":"g2"},"c":{"ida":2,"idb":2,"idc":1,"c1":"c2"}}
```

Join Predicates:

The WHERE clause must contain join predicates on all the shard key columns of all the participating tables, else the query returns an error. In the query below, the shard key of table A is not matched in the join predicate, hence the query returns an error.

```
SELECT * FROM A.B b, A a ,A.B.C c WHERE c.ida = b.ida
```

Output:

Error handling command SELECT * FROM A.B b ,A a ,A.B.C c WHERE c.ida = b.ida:
Error: at (1, 0) Tables are not joined on their common shard key

Example: Using Inner Join

Let us now see a few examples of inner join on the tables used in the [airline baggage tracking application](#).

Example 6-123 Fetch the ticket number, passenger name, and bag details of all the passengers

```
SELECT a.ticketNo, b.fullName, c.bagArrivalDate FROM ticket a,
ticket.passengerInfo b,
ticket.bagInfo c WHERE a.ticketNo = b.ticketNo AND b.ticketNo=c.ticketNo
```

Explanation: ticketpassengerInfobagInfo

Output:

```
{"ticketNo":1762324912391,"fullName":"Elane
Lemons","bagArrivalDate":"2019-03-15T10:13:00.0000Z"}
{"ticketNo":176235527825,"fullName":"Doris
Martin","bagArrivalDate":"2019-03-22T10:17:00.0000Z"}
{"ticketNo":1762344493810,"fullName":"Adam
Phillips","bagArrivalDate":"2019-02-01T16:13:00.0000Z"}
{"ticketNo":1762392135540,"fullName":"Adelaide
Willard","bagArrivalDate":"2019-02-15T21:21:00.0000Z"}
```

```
{"ticketNo":1762376407826,"fullName":"Dierdre
Amador","bagArrivalDate":"2019-03-07T13:51:00.0000Z"}
```

5 rows returned

Example 6-124 Fetch the name of the passenger whose bag's last seen station is "MEL"

```
SELECT a.fullName FROM ticket.passengerInfo a, ticket.bagInfo b WHERE
a.ticketNo = b.ticketNo AND b.lastSeenStation = "MEL"
```

Explanation: passengerInfobagInfo

Output:

```
{"fullName":"Adam Phillips"}
```

1 row returned

Example 6-125 Fetch the name of the passenger whose flight route destination is "MEL"

```
SELECT a.fullName FROM ticket.passengerInfo a, ticket.bagInfo.flightlegs b
WHERE a.ticketNo =
      b.ticketNo AND b.fltRouteDest = "MEL"
```

Explanation: passengerInfoflightlegs

Output:

```
{"fullName":"Adam Phillips"}
```

1 row returned

Inner Join vs LOJ vs NESTED TABLES

Tables in a [hierarchy](#) can be joined using the following types of joins: [Inner Join](#), [Left Outer Join](#) and [NESTED TABLES](#).

Table 6-14 Inner Join vs LOJ vs NESTED TABLES

Inner Join	Left Outer Join	NESTED TABLES
Queries multiple tables in the same hierarchy	Queries multiple tables in the same hierarchy	Queries multiple tables in the same hierarchy
Not an ANSI-SQL Standard	ANSI-SQL Standard	Not an ANSI-SQL Standard
Supports joining of tables that are in an ancestor-descendant relationship as well as tables that are not in an ancestor-descendant relationship	Supports only joining of tables that are in an ancestor-descendant relationship	Supports joining of tables that are in an ancestor-descendant relationship. Performs a union of sibling tables.

Table 6-14 (Cont.) Inner Join vs LOJ vs NESTED TABLES

Inner Join	Left Outer Join	NESTED TABLES
Returns only the matching rows	Preserves the unmatched rows from the first (left) table, joining them with a NULL row in the second (right) table	Preserves the unmatched rows from the first (left) table, joining them with a NULL row in the second (right) table
In case of 'SELECT *' where all the fields are returned, the order of fields in the result-set depends on the order of tables specified in the FROM clause.	In case of 'SELECT *' where all the fields are returned, the order of fields in the result-set is always in top-down order. That means the order of output in the result-set is always from the ancestor table first and then the descendant table. This is true irrespective of the order of tables specified in the FROM clause.	In case of 'SELECT *' where all the fields are returned, the order of fields in the result-set is always in top-down order. That means the order of output in the result-set is always from the ancestor table first and then the descendant table. This is true irrespective of the order of tables specified in the FROM clause.
Join predicate must check for matching shard keys between the participating tables	Join predicate must check for matching primary keys between the participating tables	Join predicate need not explicitly check for matching primary keys. It is implicitly done while joining the tables.

In essence, tables having ancestor-descendant relationship between them can be joined using any of the three types of join. You can choose to use one of them based on your use case. If the tables to be joined are not in an ancestor-descendant relationship, then inner join must be used.

Let us understand the 3 types of joins with a simple example shown below. Here, Table A is the parent. Table B and G are children of table A, and siblings of each other. C is the child of table B, and a descendant of table A and B. ida, idb, idc, and idg are the primary keys of tables A, B, C, and G respectively. Since A is the root table and a shard key is not explicitly defined, its primary key (ida) is taken as the shard key.

```
CREATE TABLE IF NOT EXISTS A (
    ida INTEGER,
    a1 string,
    PRIMARY KEY(ida));

CREATE TABLE IF NOT EXISTS A.B (
    idb INTEGER,
    b1 string,
    PRIMARY KEY(idb));

CREATE TABLE IF NOT EXISTS A.B.C (
    idc INTEGER,
    c1 string,
    PRIMARY KEY(idc));

CREATE TABLE IF NOT EXISTS A.G (
    idg INTEGER,
    g1 string,
    PRIMARY KEY(idg));
```


Table 6-15 Table A

ida	a1
1	a1
2	a2
3	a3

Table 6-16 Table B

(ida)	idb	b1
1	1	b1
2	2	b2

(ida) is the inherited shard key from table A.

Table 6-17 Table C

(ida)	(idb)	idc	c1
1	1	1	c1
2	2	2	c2
3	1	3	c3

(ida) is the inherited shard key from table A. (idb) is the inherited primary key from table B.

Table 6-18 Table G

(ida)	idg	g1
1	1	g1
2	2	g2

(ida) is the inherited shard key from table A.

Joining Parent-Child Tables

Inner Join:

Parent table A is joined with descendant table B. Note that only the matching rows are returned. The unmatched row from table A with ida 3, is not returned in the output.

```
SELECT * FROM A a, A.B b WHERE a.ida = b.ida ORDER BY a.ida
```

Table 6-19 Combined Table AB (every row in table A is combined with every row in table B)

ida	a1	(ida)	idb	b1
1	a1	1	1	b1
1	a1	2	2	b2
2	a2	1	1	b1

Table 6-19 (Cont.) Combined Table AB (every row in table A is combined with every row in table B)

ida	a1	(ida)	idb	b1
2	a2	2	2	b2
3	a3	1	1	b1
3	a3	2	2	b2

Table 6-20 Result after inner join (based on shard equality predicate)

ida	a1	(ida)	idb	b1
1	a1	1	1	b1
2	a2	2	2	b2

Output:

```
{ "a": { "ida": 1, "a1": "a1" }, "b": { "ida": 1, "idb": 1, "b1": "b1" } }
{ "a": { "ida": 2, "a1": "a2" }, "b": { "ida": 2, "idb": 2, "b1": "b2" } }
2 rows returned
```

Left Outer Join and NESTED TABLES:

Parent table A is joined with descendant table B. The unmatched row (ida 3) from table A is also returned with a null from table B. In this case, the behavior is the same for both left outer join and NESTED TABLES.

LOJ Query:

```
SELECT * FROM A a LEFT OUTER JOIN A.B b ON a.ida = b.ida ORDER BY a.ida
```

NESTED TABLES Query:

```
SELECT * FROM NESTED TABLES (A a descendants (A.B b)) ORDER BY a.ida
```

Table 6-21 Result after LOJ/NESTED TABLES

ida	a1	(ida)	idb	b1
1	a1	1	1	b1
2	a2	2	2	b2
3	a3	null	null	null

Output:

```
{ "a": { "ida": 1, "a1": "a1" }, "b": { "ida": 1, "idb": 1, "b1": "b1" } }
{ "a": { "ida": 2, "a1": "a2" }, "b": { "ida": 2, "idb": 2, "b1": "b2" } }
{ "a": { "ida": 3, "a1": "a3" }, "b": null }
```

Joining Sibling Tables

Inner Join:

In the below query, an inner join of tables A, B and G is performed. A is the parent, B and G are sibling tables.

```
SELECT * FROM A a, A.B b, A.G g WHERE a.ida = b.ida AND b.ida = g.ida
```

Table 6-22 Combined table AB (every row in table A is combined with every row in table B)

ida	a1	(ida)	idb	b1
1	a1	1	1	b1
1	a1	2	2	b2
2	a2	1	1	b1
2	a2	2	2	b2
3	a3	1	1	b1
3	a3	2	2	b2

Table 6-23 Combined table ABG (every row in table AB is combined with every row in table G)

ida	a1	(ida)	idb	b1	(ida)	idg	g1
1	a1	1	1	b1	1	1	g1
1	a1	1	1	b1	2	2	g2
1	a1	2	2	b2	1	1	g1
1	a1	2	2	b2	2	2	g2
2	a2	1	1	b1	1	1	g1
2	a2	1	1	b1	2	2	g2
2	a2	2	2	b2	1	1	g1
2	a2	2	2	b2	2	2	g2
3	a3	1	1	b1	1	1	g1
3	a3	1	1	b1	2	2	g2
3	a3	2	2	b2	1	1	g1
3	a3	2	2	b2	2	2	g2

Table 6-24 Result after inner join (based on shard key equality predicate)

ida	a1	(ida)	idb	b1	(ida)	idg	g1
1	a1	1	b1	1	1	1	g1
2	a2	2	b2	2	2	2	g2

Output:

```
{"a":{"ida":1,"a1":"a1"},"b":{"ida":1,"idb":1,"b1":"b1"},"g":
{"ida":1,"idg":1,"g1":"g1"}}
{"a":{"ida":2,"a1":"a2"},"b":{"ida":2,"idb":2,"b1":"b2"},"g":
```

```
{"ida":2,"idg":2,"g1":"g2"}}
```

2 rows returned

Left Outer Join:

Sibling join is not possible.

```
SELECT * FROM A a LEFT OUTER JOIN A.B b ON a.ida=b.ida LEFT OUTER JOIN A.G g
ON a.ida = g.ida
```

Output:

Error handling command `SELECT * FROM A a LEFT OUTER JOIN A.B b ON a.ida=b.ida LEFT OUTER JOIN A.G g ON a.ida = g.ida`: Error: at (1, 75) Table A.G is not descendant of table A.B. Tables in left-outer-joins must appear in top-down order after the target table

NESTED TABLES:

```
SELECT * FROM NESTED TABLES (A a descendants (A.b b, A.G g))
```

Union of sibling tables B and G is performed

Table 6-25 Union of sibling tables B and G

(ida)	idb	b1	(ida)	idg	g1
1	1	b1	null	null	null
2	2	b2	null	null	null
null	null	null	1	1	g1
null	null	null	2	2	g2

Table 6-26 Result after NESTED TABLES join (LOJ of A with union of B and G)

ida	a1	(ida)	idb	b1	(ida)	idg	g1
1	a1	1	1	b1	null	null	null
1	a1	null	null	null	1	1	g1
2	a2	2	2	b2	null	null	null
2	a2	null	null	null	2	2	g2
3	a3	null	null	null	null	null	null

Output:

```
{"a":{"ida":1,"a1":"a1"},"b":{"ida":1,"idb":1,"b1":"b1"},"g":null}
{"a":{"ida":1,"a1":"a1"},"b":null,"g":{"ida":1,"idg":1,"g1":"g1"}}
{"a":{"ida":2,"a1":"a2"},"b":{"ida":2,"idb":2,"b1":"b2"},"g":null}
{"a":{"ida":2,"a1":"a2"},"b":null,"g":{"ida":2,"idg":2,"g1":"g2"}}
{"a":{"ida":3,"a1":"a3"},"b":null,"g":null}
```

5 rows returned

Joining Tables that are not in an Ancestor-Descendant Relationship**Inner Join:**

Table G is joined with C. Parent of table G is A and parent of table C is B. Table G and C are not in an ancestor-descendant relationship.

```
SELECT * FROM A.G g, A.B.C c WHERE g.ida=c.ida
```

Table 6-27 Combined Table GC (every row in table G is combined with every row in table C)

(ida)	idg	g1	(ida)	(idb)	idc	c1
1	1	g1	1	1	1	c1
1	1	g1	2	2	2	c2
1	1	g1	3	1	3	c3
2	2	g2	1	1	1	c1
2	2	g2	2	2	2	c2
2	2	g2	3	1	3	c3

Table 6-28 Result after inner join (based on shard key equality predicate)

(ida)	idg	g1	(ida)	(idb)	idc	c1
1	1	g1	1	1	1	c1
2	2	g2	2	2	2	c2

NESTED TABLES and LOJ cannot be used to join tables that do not have an ancestor-descendant relationship between them.

7

Data Row Management

This chapter describes data rows and inserting and managing data rows in Oracle NoSQL Database.

This chapter contains the following topics:

- [INSERT Statement](#)
- [Inserting Rows into JSON Collection Tables](#)
- [Inserting Rows with an IDENTITY Column](#)
- [Inserting rows into a table with a UUID column](#)
- [Inserting rows with an MR_COUNTER column](#)
- [Upsert statement](#)
- [Updating rows of a table with a UUID column](#)
- [DELETE Statement](#)
- [UPDATE Statement](#)
 - [Update Clauses](#)
 - * [SET Clause](#)
 - * [ADD Clause](#)
 - * [PUT Clause](#)
 - * [REMOVE Clause](#)
 - * [SET TTL Clause](#)
 - [Updating rows with an IDENTITY Column](#)
 - [Updating rows with an MR_COUNTER column](#)
 - [Example: Updating Rows](#)
 - [Example: Updating JSON Data](#)
 - [Example: Updating JSON collection tables](#)
 - [Example: Updating TTL](#)
 - [Example: Updating IDENTITY defined as GENERATED ALWAYS](#)
 - [Example: Updating IDENTITY defined as GENERATED BY DEFAULT](#)
- [JSON Collection Table Example](#)

INSERT Statement

The INSERT statement is used to construct a new row and add it in a specified table.

Syntax

```
insert_statement ::=
    [variable_declaration]
    INSERT INTO table_name
    [[AS] table_alias]
    ["(" id ("," id)* ")"]
    VALUES "(" insert_clause ("," insert_clause)* ")"
    [SET TTL ttl_clause ]
    [returning_clause]
```

```
insert_clause ::= DEFAULT | expression
```

```
returning_clause ::= RETURNING select_list
```

Semantics

The row will be inserted only if it does not exist already.

Insert statements may start with declarations of external variables that are used in the rest of the statement. See [Variable Declaration](#). However, contrary to queries, external variables can be used in inserts without being declared. This is because inserts do not query any data, and as result, knowing the type of external variables in advance is not important as there isn't any query optimization to be done.

Optional column(s) may be specified after the table name. This list contains the column names for a subset of the table's columns. The subset must include all the primary key columns. If no columns list is present, the default columns list is the one containing all the columns of the table, in the order they are specified in the [CREATE TABLE Statement](#) section.

The columns in the columns list correspond one-to-one to the expressions (or DEFAULT keywords) listed after the VALUES clause (an error is raised if the number of expressions/DEFAULTs is not the same as the number of columns). These expressions/DEFAULTs compute the value for their associated column in the new row. Specifically, each expression is evaluated and its returned value is cast to the type of its associated column. The cast behaves like the cast expression as described in the [Cast Expression](#) section. An error is raised if an expression returns more than one item. If an expression returns no result, NULL is used as the result of that expression. If instead of an expression, the DEFAULT keyword appears in the VALUES list, the default value of the associated column is used as the value of that column in the new row. The default value is also used for any missing columns, when the number of columns in the columns list is less than the total number of columns in the table.

The expressions in the VALUES list may reference external variables, which unlike query statements, do not need to be declared in a declarations section.

Following the VALUES list a SET TTL clause may be used to set the expiration time of the new row. As described in the [CREATE TABLE Statement](#) section, every row has an expiration time, which may be infinite, and which is computed in terms of a Time-To-Live (TTL) value that is specified as a number of days or hours. Specifically, for a TTL value of N hours/days, where N is greater than zero, the expiration time is computed as the current time (in UTC) plus N hours/days, rounded up to the next full hour/day. For example, if the current time is 2017-06-01T10:05:30.0 and N is 3 hours, the expiration time will be 2017-06-01T14:00:00.0. If N is 0, the expiration time is infinite.

As shown in the syntax, the SET TTL clause comes in two flavors. When the USING TABLE DEFAULT syntax is used, the TTL value is set to the table default TTL that was specified in the CREATE TABLE statement. Otherwise, the SET TTL contains an expression, which computes

a new TTL value. If the result of this expression is empty, the default TTL of the table is used. Otherwise, the expression must return a single numeric item, which is cast to an integer N. If N is negative, it is set to 0. To the right of the TTL expression, the keyword HOURS or DAYS must be used to specify whether N is a number of hours or days, respectively.

If the insert statement contains a SET TTL clause, an expiration time is computed as described above and applied to the row being inserted. If no SET TTL clause is used, the default table TTL is used to compute the expiration time of the inserted row. In case of MR Tables, when this row is replicated to other regions, its expiration time is also replicated as an absolute timestamp value. Therefore, the replicated rows will expire along with the original row, irrespective of when they were replicated. If the same row is inserted with a TTL value in multiple regions, then the TTL value will be set in all regions to the value held in the row with the greatest write timestamp.

The last part of the insert statement is the RETURNING clause. If not present, the result of the update statement is a record with a single field whose name is "NumRowsInserted" and whose value is the number of rows inserted: 0 if the row existed already, or 1 otherwise. Otherwise, if there is a RETURNING clause, it acts the same way as the SELECT clause: it can be a "*", in which case, a full row will be returned, or it can have a list of expressions specifying what needs to be returned. In the case of an INSERT where no insertion is actually done (because the row exists already), the RETURNING clause acts on the existing row. Otherwise, the RETURNING clause acts on the new row.

Using DEFAULT values while inserting data

The following statement inserts a row to the users table from [CREATE TABLE Statement](#). Notice that the value for the expenses column will be set to NULL, because the DEFAULT clause is used for that column.

Example 7-1 Insert a row to the users table using DEFAULT

```
INSERT INTO users VALUES (
    10,
    "John",
    "Smith",
    [ {"first" : "Johny", "last" : "BeGood"} ],
    22,
    45000,
    { "street" : "Main", "number" : 10, "city" : "Reno", "state" : "NV"},
    [ 30, 55, 43 ],
    DEFAULT
)
```

Using a RETURNING clause in an INSERT statement.

A RETURNING clause acts the same way as the SELECT clause. A "*", in the example below returns the full row that got inserted.

Example 7-2 Insert a row to the users table using RETURNING clause

```
INSERT INTO users VALUES (
    20,
    "Mary",
    "Ann",
    null,
    43,
    90000,
```



```

        { "street" : "Main", "number" : 89, "city" : "San Jose", "state" : "CA"},
        null,
        DEFAULT
    ) RETURNING *

```

Output:

```

{"id":20,"firstName":"Mary","lastName":"Ann","otherNames":null,
"age":43,"income":90000,"address":{"city":"San Jose","number":89,
"state":"CA","street":"Main"},"connections":null,"expenses":null}

```

Setting a TTL for a row while inserting a row

In the example below, the expiration of the row that is inserted is set to 2 days.

Example 7-3 Insert a row to the users table by using a SET clause to set the TTL value

```

INSERT INTO users VALUES (
    30,
    "Peter",
    "Paul",
    null,
    25,
    53000,
    { "street" : "Main", "number" : 3, "city" : "Fresno", "state" : "CA"},
    null,
    DEFAULT
)SET TTL 2 days

```

Output:

```

{"NumRowsInserted":1}

```

Inserting data into a child table:

A child table inherits the primary key columns of its parent table. This is done implicitly, without including the parent columns in the CREATE TABLE statement of the child. So the child table has additional columns pertaining to the primary key of the parent table. While inserting data into a child table, the value for the primary key of the parent table should also be inserted.

Example 7-4 Insert a row to the userdet child table

The description of the child table userdet is as follows:

```
desc as json table users.userdet
```

Output:

```

{
  "json_version" : 2,
  "type" : "table",
  "name" : "userdet",
  "id" : "63",
  "parent" : "users",

```

```

"regions" : {
  "2" : "FRA",
  "1" : "LON"
},
"fields" : [{
  "name" : "id",
  "type" : "INTEGER",
  "nullable" : false
}, {
  "name" : "pan",
  "type" : "INTEGER",
  "nullable" : false
}, {
  "name" : "address",
  "type" : "STRING",
  "nullable" : true
}, {
  "name" : "email",
  "type" : "STRING",
  "nullable" : true
}],
"primaryKey" : ["id", "pan"],
"shardKey" : ["id"]
}

```

When you retrieve data from a child table, the primary key of the parent table is also retrieved.

```
SELECT * FROM users.userdet
```

Output:

```
{"id":1,"pan":100,"address":"bangalore","email":"test@one.com"}
```

Here "id" is the primary key of the parent table. While inserting data, provide the value for the "id" column. Else, an error is returned.

```
insert into users.userdet values(1,100,"bangalore","test@one.com")
```

Output:

```
{"NumRowsInserted":1}
```

Inserting Rows into JSON Collection Tables

JSON collection tables provide the flexibility to declare [primary key fields](#) during the table creation. The value of primary key fields along with the other fields in the document must be supplied during the insertion of data.

When you insert data, each row is inserted as a single document containing any number of JSON fields. To insert data, you must either supply the values for the fields inside a JSON object ({}), or state them explicitly.

Use one of the following methods to insert the data into a JSON collection table:

- **Using explicitly declared field names:** You must explicitly supply the primary key field followed by the top-level JSON field names in the INSERT statement. You include the corresponding values using the **values** clause. The first value must be the value associated with the primary key field. This must be followed by the values associated with the corresponding field names supplied for the JSON document.

Example 7-5 Insert data into the JSON collection table created for a shopping application through explicit declaration of field names

```
INSERT into storeAcct(contactPhone, firstName, lastName, address, cart)
values("1817113382", "Adam", "Smith", {"street" : "Tex Ave", "number" : 401,
"city" : "Houston", "state" : "TX", "zip" : 95085}, [{"item" : "handbag",
"quantity" : 1, "priceperunit" : 350}, {"item" : "Lego", "quantity" : 1,
"priceperunit" : 5500}])
```

Explanation: In the above example, you insert the shopper's data into the [shopping application table](#) by supplying the contactPhone as the primary key followed by other details of the shopper as JSON fields. The shopper's details are stored internally as JSON data. Notice that in JSON collection tables, you do not supply a column name for the document itself, only provide the JSON fields in the document.

You can add another row to the same table with additional fields.

```
INSERT into storeAcct(contactPhone, firstName, lastName, gender, address,
notify, cart, wishlist) values("1917113999", "Sharon", "Willard", "F",
{"street" : "Maine", "number" : 501, "city" : "San Jose", "state" : "San
Francisco", "zip" : 95095}, "yes", [{"item" : "wallet", "quantity" : 2,
"priceperunit" : 950}, {"item" : "wall art", "quantity" : 1, "priceperunit" :
9500}], [{"item" : "Tshirt", "priceperunit" : 500}, {"item" : "Jenga",
"priceperunit" : 850}])
```

In the above statement, you insert the shopper's data with additional fields such as gender, notify, and wishlist as compared with the first inserted row. The wishlist field is a JSON array that includes the details of items wishlisted by the shopper.

- **Using positional values:** You must supply the primary key field values followed by document fields as name/value pairs encapsulated in a single JSON object. The fields in the document must adhere to [JSON data type](#) format. Any non-conformance will lead to exceptions and display error messages.

Example 7-6 Insert data into the JSON collection table created for a shopping application by supplying positional values

```
INSERT into storeAcct values("1817113382", {"firstName" : "Adam",
"lastName" : "Smith", "address" : {"street" : "Tex Ave", "number" : 401,
"city" : "Houston", "state" : "TX", "zip" : 95085}, "cart" : [{"item" :
"handbag", "quantity" : 1, "priceperunit" : 350}, {"item" : "Lego",
"quantity" : 1, "priceperunit" : 5500}]})
```

Explanation: In this insert statement, you insert the same data from [example 1](#) above using the positional value of the elements. You use the **values** clause followed by the value for the primary key field and then a JSON object {} containing the rest of the name/value pairs for the document.

Additional examples:

Example 7-7 Insert data into the JSON collection table for a shopping application with an MR_COUNTER

```
INSERT into storeAcctMR(contactPhone, firstName, lastName, address, cart,
mycounter) values("1817113382", "Adam", "Smith", {"street" : "Tex Ave",
"number" : 401, "city" : "Houston", "state" : "TX", "zip" : 95085},
[{"item" : "handbag", "quantity" : 1, "priceperunit" : 350}, {"item" : "Lego",
"quantity" : 1, "priceperunit" : 5500}], 0)
```

Explanation: While inserting data into the [JSON collection table with an MR_COUNTER](#), you must supply a value of 0 for the [MR_COUNTER](#).

Note

- If you supply a non-zero value, the system will ignore the supplied value and insert the value 0.
- You can't supply the keyword DEFAULT in the INSERT clause for MR_COUNTER value.
- The system returns an error if the MR_COUNTER value is skipped or DEFAULT is supplied in the INSERT statement. For more details, see [Inserting rows into a JSON column having MR_COUNTER data type](#).

```
SELECT contactPhone, firstName, mycounter FROM storeAcctMR where contactPhone
= "1817113382"
```

Output:

```
{
  "contactPhone" : "1817113382",
  "firstName" : "Adam",
  "mycounter" : 0
}
```

Example 7-8 Insert shopper's data with more than one primary key field in the JSON collection table

```
INSERT into storeAcctComposite(contactPhone, firstName, lastName, address,
orders) values("1417114488", "Doris", "Martin", {"Dropbox" : "Presidency
College", "city" : "Kansas City", "state" : "Alabama", "zip" : 95065})
RETURNING *
```

Explanation: The storeAcctComposite table includes two primary key fields, the contactPhone field and the id field. For table creation, see [JSON collection composite keys](#). As the order IDs for the purchased items are autogenerated in the shopping application, the id field is declared as an IDENTITY column while creating the table. While inserting data, you must supply the contactPhone primary key value along with document fields. If you do not supply the id field value, the system autogenerates the IDENTITY field value. For more details on the IDENTITY column, see [Inserting Rows with an IDENTITY Column](#).

Output:

```
{
  "contactPhone" : "1417114488",
  "id" : 1,
  "address" : {
    "Dropbox" : "Presidency College",
    "city" : "Kansas City",
    "state" : "Alabama",
    "zip" : 95065
  },
  "firstName" : "Doris",
  "lastName" : "Martin"
}
```

Using APIs to Insert Data into JSON Collection Tables

You can create a JSON collection table and insert data into it using supported language drivers (SDKs). You can define the primary key fields using an allowed subset of NoSQL data types. For more information, see [Supported data type](#).

The language SDKs handle the creation of JSON collection tables or the insertion of data in the same way as fixed schema tables. For examples on the creation of JSON collection tables using APIs and inserting data into the tables, see *Using APIs to create tables* in Creating a table section of the *Developer's Guide*.

Inserting Rows with an IDENTITY Column

The system generates an IDENTITY column value when the keyword DEFAULT is used as the insert_clause for the IDENTITY column. Here are a few examples that show INSERT statements for both flavors of the IDENTITY column – GENERATED BY DEFAULT and GENERATED ALWAYS.

How you create an IDENTITY field affects what happens when you INSERT values. You cannot change the IDENTITY value of a column that is a primary key.

Example 7-9 Create a table with an IDENTITY field as GENERATED ALWAYS Where the IDENTITY field is not a primary key

```
CREATE Table Test_SGSqlInsert2(
  id INTEGER,
  name STRING,
  deptId INTEGER GENERATED ALWAYS AS IDENTITY (CACHE 1),
  PRIMARY KEY(id))

INSERT INTO Test_SGSqlInsert2 VALUES (148, 'sally', DEFAULT)
INSERT INTO Test_SGSqlInsert2 VALUES (250, 'joe', DEFAULT)
INSERT INTO Test_SGSqlInsert2 VALUES (346, 'dave', DEFAULT)
```

The above INSERT statement will insert the following rows. The system generates values 1, 2, and 3 for IDENTITY column deptId.

```
148, 'sally', 1
250, 'joe', 2
346, 'dave', 3
```

To get the value of the generated deptId for future reference in one statement, use the returning deptId clause as follows:

```
INSERT INTO Test_SGSqlInsert2 VALUES (600, 'jabba', DEFAULT) returning deptId
6
INSERT INTO Test_SGSqlInsert2 VALUES (700, 'bubba', DEFAULT) returning deptId
7
```

Using the following INSERT statement causes an exception since the user supplied a value of 200 for an IDENTITY GENERATED ALWAYS column. You cannot specify any value for any IDENTITY field you define as GENERATED BY DEFAULT AS IDENTITY.

```
INSERT INTO Test_SGSqlInsert2 VALUES (1, 'joe', 200)
```

Output:

```
Error handling command INSERT INTO Test_SGSqlInsert2 VALUES (1, 'joe', 200):
Error: at (1, 48) Generated always identity column must use DEFAULT construct.
```

Example 7-10 Create a table with a DeptID integer field, GENERATED BY DEFAULT AS IDENTITY, and make it the primary and shard key

```
CREATE TABLE Test_SGSqlInsert_Default (
  ID INTEGER,
  NAME STRING,
  DeptID INTEGER GENERATED BY DEFAULT AS IDENTITY (
    START WITH 1
    INCREMENT BY 1
    MAXVALUE 100),
  PRIMARY KEY (SHARD(DeptID), ID))
```

The following statements show how to insert values into table Test_SGSqlInsert_Default. In this case, since the column ID is not an IDENTITY, you can assign integer values to the field:

```
INSERT INTO Test_SGSqlInsert_Default VALUES (100, 'tim', DEFAULT)
INSERT INTO Test_SGSqlInsert_Default VALUES (200, 'dave', 210)
INSERT INTO Test_SGSqlInsert_Default VALUES (300, 'sam', 310)
INSERT INTO Test_SGSqlInsert_Default VALUES (400, 'Jennifer', DEFAULT)
INSERT INTO Test_SGSqlInsert_Default VALUES (500, 'Barbara', 2)
```

These sample statements insert the following rows into the database.

```
300, 'sam', 310
100, 'tim', 1
400, 'Jennifer', 2
```

```
500, 'Barbara', 2
200, 'dave', 210
```

Since you specified two values as DEFAULT in your INSERT statements, the SG supplies them, as 1 and 2. The other values are inserted as you specify (210, 310, and 2). Each value is acceptable, even though one results in two DeptID values the same (2 supplied from a DEFAULT, and 2 as a value you supply).

Because you defined the IDENTITY column as GENERATED BY DEFAULT AS IDENTITY, the SG supplies a value only when you do not specify a value. Specifying values 210, 310, or 2 is correct. The system neither checks for duplicates, nor enforces uniqueness for GENERATED BY DEFAULT AS IDENTITY column values. It is the application's responsibility to ensure that there are no duplicate values if that is a requirement.

Example 7-11 Inserting rows into a table with primary key as IDENTITY Column GENERATED BY DEFAULT

```
CREATE TABLE Test_SGSqlInsert_Default_id (
    id INTEGER GENERATED BY DEFAULT AS IDENTITY,
    name STRING,
    deptId INTEGER,
    PRIMARY KEY (id))

INSERT INTO Test_SGSqlInsert_Default_id VALUES (100, 'tim', 3)
INSERT INTO Test_SGSqlInsert_Default_id VALUES (DEFAULT, 'dave', 210)
INSERT INTO Test_SGSqlInsert_Default_id VALUES (300, 'sam', 310)
INSERT INTO Test_SGSqlInsert_Default_id VALUES (500, 'Jennifer', 410)
INSERT INTO Test_SGSqlInsert_Default_id (name,deptId) VALUES ('Abby', 510)
```

The above statements will insert the following rows into the database. Notice that when you supply DEFAULT for the id field, the SG auto-generates an id value because the primary key field (id) is defined as the IDENTITY column. Similarly, in the last insert statement where you supply only the name and deptID field values, the SG auto-generates the id value.

```
{"id":2,"name":"Abby","deptId":4}
{"id":300,"name":"sam","deptId":310}
{"id":100,"name":"tim","deptId":3}
{"id":1,"name":"dave","deptId":210}
{"id":500,"name":"Jennifer","deptId":410}
```

Inserting rows into a table with a UUID column

The system generates a UUID column value when the keyword DEFAULT is used as the insert_clause for the UUID column.

Here are a few examples that show INSERT statements for both flavors of the UUID column – GENERATED BY DEFAULT and when no DEFAULT CLAUSE is specified in a CREATE TABLE statement. The keyword DEFAULT in the INSERT statement applies only when the UUID column is declared as GENERATED BY DEFAULT.

Example 7-12 Inserting rows into a table with a UUID column without GENERATED BY DEFAULT clause

```
CREATE TABLE myTable (id STRING AS UUID, name STRING, PRIMARY KEY (id))
```

Output:

Statement completed successfully

```
INSERT INTO myTable values("a81bc81b-dead-4e5d-abff-90865d1e13b1","test1")
```

Output:

Statement completed successfully

In the above example, the id column in the table myTable has no "GENERATED BY DEFAULT" defined. Therefore, whenever you insert a new row, you need to explicitly specify the value for the id column.

Example 7-13 Inserting rows into a table with a UUID column using the random_uuid function

The value for a UUID column can also be generated using the random_uuid function. See Function to generate a UUID string.

```
INSERT INTO myTable values(random_uuid(),"test2")
```

Output:

```
{"NumRowsInserted":1}  
1 row returned  
Statement completed successfully
```

```
select * from myTable
```

Output:

```
{"id":"d576ab3b-8a36-4dff-b50c-9d9d4ca6072c","name":"test2"}  
{"id":"a81bc81b-dead-4e5d-abff-90865d1e13b1","name":"test1"}  
2 rows returned  
Statement completed successfully
```

In this example, a randomly generated UUID is fetched using the [random_uuid\(\)](#) function. This value is used in the INSERT statement.

Example 7-14 Inserting rows into a table with a UUID column with GENERATED BY DEFAULT clause

```
CREATE TABLE myTable (id STRING AS UUID GENERATED BY DEFAULT, name STRING,  
PRIMARY KEY (id))
```


Output:

Statement completed successfully

```
INSERT INTO myTable VALUES(default,"test1") returning id
```

Output:

```
{"id":"e7fbab63-7730-4ec9-be73-a62e33ea73c3"}  
Statement completed successfully
```

In the above example, the id column in myTable has "GENERATED BY DEFAULT" defined. The system generates a UUID column value when the keyword DEFAULT is used in the insert_clause for the UUID column. The system generated UUID value is fetched using the returning clause.

Inserting rows with an MR_COUNTER column

While data is inserted in a multi-region table with an MR_COUNTER column, the system generates a default value of zero for the MR_COUNTER column value in the following two cases.

- When the keyword DEFAULT is used in the INSERT clause for the MR_COUNTER column.
- When the MR_COUNTER column is skipped in the INSERT clause.

Specifying DEFAULT clause while inserting data into an MR_COUNTER column**Example 7-15 Insert a row to a multi-region table using the DEFAULT clause**

Create a multi-region table in DEN and LON regions with an MR_COUNTER:

```
CREATE Table myTable( name STRING,  
                      count INTEGER AS MR_COUNTER,  
                      PRIMARY KEY(name)) IN REGIONS DEN,LON
```

Insert a row to the table by specifying the DEFAULT clause to the counter:

```
INSERT INTO myTable VALUES ("Bob", DEFAULT)
```

Fetch the counter value:

```
SELECT * FROM myTable
```

Output:

```
{"name":"Bob", "count":0}
```

Skipping the MR_COUNTER column while inserting data into a multi-region table**Example 7-16 Insert a row to a multi-region table without specifying a value to the MR_COUNTER**

Create a multi-region table in DEN and LON regions with an MR_COUNTER:

```
CREATE Table myTable( name STRING,  
                      count INTEGER AS MR_COUNTER,  
                      PRIMARY KEY(name)) IN REGIONS DEN,LON
```

Insert a row to the table by specifying the column names, without specifying the counter or supplying any value to it:

```
INSERT INTO myTable(name) VALUES ('Chris')
```

Fetch the table data. Notice that the system automatically inserts a zero value to the counter.

```
SELECT * FROM myTable
```

Output:

```
{"name": "Chris", "count": 0}
```

Error when MR_COUNTER column is not skipped and no DEFAULT clause is given

While inserting all column values to a row, you must specify the DEFAULT clause for the MR_COUNTER column. Else, the system displays an error.

Example 7-17 Insert all columns in a row without specifying a value to the MR_COUNTER

Create a multi-region table in DEN and LON regions with an MR_COUNTER:

```
CREATE Table myTable( name STRING,  
                      count INTEGER AS MR_COUNTER,  
                      PRIMARY KEY(name)) IN REGIONS DEN,LON
```

Insert all the columns in a row. Notice that you don't supply a DEFAULT value to the counter.

```
INSERT INTO myTable VALUES ('Chris')
```

Output:

```
Error handling command execute 'INSERT INTO myTable VALUES ('Chris')':  
Error: at (1, 0) The number of VALUES expressions is not equal to the number  
of table columns
```

You cannot insert values into an MR_COUNTER column explicitly. A DEFAULT construct must always be used or the MR_COUNTER column should be skipped in the INSERT clause. If you try to insert values into the MR_COUNTER column using the INSERT clause or using API, an error is thrown as shown below.

Example 7-18 Insert a row to a multi-region table by specifying a value other than DEFAULT to the MR_COUNTER

```
CREATE Table myTable ( name STRING,  
                      count INTEGER AS MR_COUNTER,  
                      PRIMARY KEY(name)) IN REGIONS DEN
```

```
INSERT INTO myTable VALUES("Tom",0)'
```

Output:

*Error handling command execute 'INSERT INTO myTable VALUES("Tom",0)': Error:
at (1, 38) MRCounter column must use DEFAULT construct.*

Inserting rows into a JSON column having MR_COUNTER data type:

When inserting a row into the multi-region table with a JSON MR_COUNTER, you must supply a value zero to the MR_COUNTER.

Note

- The system initially assigns a value of zero to all MR_COUNTER data types even if you explicitly supply a non-zero value. This also holds good when you try to provide a value that is not an INTEGER or LONG or NUMBER.
- You can't supply the keyword DEFAULT while inserting a JSON MR_COUNTER.
- The system will return an error if you try to insert data into an MR table without supplying a value to the declared JSON MR_COUNTER field or using the keyword DEFAULT.

Example 7-19 Insert data into multi-region table with a JSON MR_COUNTER

```
INSERT INTO demoJSONMR VALUES ("Anna",  
    {  
        "id" : 1,  
        "counter" : 0,  
        "person" : {  
            "age" : 10,  
            "count" : 0,  
            "number" : 100  
        }  
    }  
)
```

```
SELECT * FROM demoJSONMR
```

Output:

```
{"name":"Anna","jsonWithCounter":{"id" : 1,"counter":0,  
                                "person":{"age":10,"count":0,"number":100}}
```

```
}  
}
```

Upsert statement

The word UPSERT combines UPDATE and INSERT, describing the statement's function.

Syntax:

```
upsert_statement ::=  
    [variable_declaration]  
    UPSERT INTO table_name  
    [[AS] table_alias]  
    ["(" id ("," id)* ")"]  
    VALUES "(" insert_clause ("," insert_clause)* ")"  
    [SET TTL ttl_clause ]  
    [returning_clause]
```

```
insert_clause ::= DEFAULT | expression
```

```
returning_clause ::= RETURNING select_list
```

Use an UPSERT statement to insert a row where it does not exist, or to update the row with new values when it does.

- Optional column(s) may be specified after the table name. This list contains the column names for a subset of the table's columns. The subset must include all the primary key columns. If no columns list is present, the default columns list is the one containing all the columns of the table, in the order they are specified in the [CREATE TABLE Statement](#).
- The columns in the columns list correspond one-to-one to the expressions (or DEFAULT keywords) listed after the VALUES clause (an error is raised if the number of expressions/DEFAULTs is not the same as the number of columns).
- Following the VALUES list a SET TTL clause may be used to set the expiration time of an upserted(inserted/updated) row.
- If there is no RETURNING clause, the result of the UPSERT statement is a record with a single field whose name is **NumRowsInserted** and whose value is the number of rows inserted: 0 if the row existed already and an update happened, or 1 otherwise. If there is a RETURNING clause, it acts the same way as the SELECT clause: it can be a *, in which case, a full row will be returned, or it can have a list of expressions specifying what needs to be returned.

The users table has three rows as shown below:

```
SELECT count(*) FROM users
```

Output:

```
{"Column_1":3}
```

Example 7-20 Update data in the users table using UPSERT command

The existing value for an user with id 10 is shown below.

```
SELECT * FROM users WHERE id=10
```

Output

```
{
  "id" : 10,
  "firstName" : "John",
  "lastName" : "Smith",
  "otherNames" : [{
    "first" : "Johny",
    "last" : "BeGood"
  }],
  "age" : 22,
  "income" : 45000,
  "address" : {
    "city" : "Reno",
    "number" : 10,
    "state" : "NV",
    "street" : "Main"
  },
  "connections" : [30, 55, 43],
  "expenses" : null
}
```

You modify the existing row using the UPSERT command. The otherNames array and income field is modified.

```
UPSERT INTO users VALUES (
  10,
  "John",
  "Smith",
  [ {"first" : "Johny", "last" : "AlwaysGood"} ],
  22,
  80000,
  { "street" : "Main", "number" : 10, "city" : "Reno", "state" : "NV"},
  [ 30, 55, 43 ],
  DEFAULT
)
```

Output:

```
{"NumRowsInserted":0}
```

1 row returned

You get the result as {"NumRowsInserted":0} which implies that the row has been updated. The updated value for an user with id 10 can be verified with a SELECT statement as shown below.

```
SELECT * FROM users WHERE id=10
```

Output:

```
{
  "id" : 10,
  "firstName" : "John",
  "lastName" : "Smith",
  "otherNames" : [{
    "first" : "Johnny",
    "last" : "AlwaysGood"
  }],
  "age" : 22,
  "income" : 80000,
  "address" : {
    "city" : "Reno",
    "number" : 10,
    "state" : "NV",
    "street" : "Main"
  },
  "connections" : [30, 55, 43],
  "expenses" : null
}
```

Example 7-21 Update only some columns using UPSERT statement

If you use the UPSERT statement and specify the values of only few columns in the VALUES clause but do not specify the corresponding column names in the INTO clause, you get an error stating the number of columns in the table do not match with the number of values in the VALUES clause as shown below.

```
UPSERT INTO users VALUES ( 11, "John", "Smith" )
```

Output:

```
Error handling command UPSERT INTO users VALUES ( 11, "John", "Smith" ):
Error: at (1, 0) The number of VALUES expressions is not equal to the number
of table columns
```

You can avoid this error by specifying the column list after the table name. Here if you do not supply values for all the columns in a UPSERT statement, then those columns get a DEFAULT value if such an option is specified in the corresponding CREATE TABLE statement or those columns are assigned NULL values as shown below.

```
UPSERT INTO users(id,firstName,lastName) VALUES (11,"John","Smith")
```

Output:

```
{"NumRowsInserted":1}  
1 row returned
```

You get the result as {"NumRowsInserted":1} which implies that one new row has been inserted. The updated value for an user with id 11 can be verified with a SELECT statement as shown below.

```
SELECT * FROM users WHERE id=11
```

Output:

```
{  
  "id" : 11,  
  "firstName" : "John",  
  "lastName" : "Smith",  
  "otherNames" : null,  
  "age" : null,  
  "income" : null,  
  "address" : null,  
  "connections" : null,  
  "expenses" : null  
}
```

All the fields which were not part of the UPSERT statement has a NULL value.

Note

Even if a column has a non NULL value (for example lastNames in the query above), it can becomes NULL if it is omitted in a subsequent UPSERT statement as shown below. Here you are using an optional RETURNING statement to fetch the data after the UPSERT is performed.

```
UPSERT INTO users(id,firstName) VALUES (11,"Joseph") returning *
```

Output:

```
{  
  "id" : 11,  
  "firstName" : "Joseph",  
  "lastName" : null,  
  "otherNames" : null,  
  "age" : null,  
  "income" : null,  
  "address" : null,  
  "connections" : null,  
  "expenses" : null  
}
```

The column lastNames has become NULL because of the UPSERT statement.

Example 7-22 Add a new shopper's record to the storeAcct table.

You can use the UPSERT statement to add a new document or update fields in an existing document in a JSON collection table. Consider the JSON collection table created for a [shopping application](#).

```
UPSERT into storeAcct values ("1417114588", {"firstName" : "Dori",  
"lastName" : "Martin", "email" : "dormartin@usmail.com", "address" :  
{"Dropbox" : "Presidency College"}}) RETURNING *
```

Explanation: In the above example, you use the UPSERT statement to add a new row to the storeAcct table.

You can use the UPSERT statement to update a shopper's information. Only the fields supplied in the UPSERT statement are updated in the document. The omitted fields are removed from the document.

Output:

```
{"contactPhone":"1417114588","address":{"Dropbox":"Presidency  
College"},"email":"lorphil@usmail.com","firstName":"Dori","lastName":"Martin"}
```

Updating rows of a table with a UUID column

You can update a UUID column whether or not it is GENERATED BY DEFAULT. You can use the function random_uuid to generate a random UUID value to update the column. The function [random_uuid](#) returns a randomly generated UUID, as a string of 36 characters.

Example 7-23 Updating a UUID Column defined without GENERATED BY DEFAULT clause

Create a table with id column as a UUID data type.

```
CREATE TABLE myTable (tabId INTEGER, id STRING AS UUID, PRIMARY KEY (tabId))
```

Output:

Statement completed successfully

Insert data into the table:

```
INSERT INTO myTable values(1, "a81bc81b-dead-4e5d-abff-90865d1e13b1")
```

Output:

Statement completed successfully

Update the id column using the built-in random_uuid function.

```
UPDATE myTable set id=random_uuid() where tabId=1
```


Output:

Statement completed successfully

The example above shows how you can update a UUID column which is NOT GENERATED BY DEFAULT. To do so, the UUID column should not be part of the primary key, as NoSQL Primary key values are immutable. In the above example, `tabId` is the Primary key. So you can update the UUID column using the `random_uuid` function.

DELETE Statement

The DELETE statement is used to remove from a table a set of rows satisfying a condition.

Syntax

```
delete_statement ::=
    [variable_declaration]
    DELETE FROM table_name [[AS] table_alias]
    WHERE expression
    [returning_clause]

returning_clause ::= RETURNING select_list
```

Semantics

The delete statement is used to delete from a table a set of rows satisfying a condition. The condition is specified in a WHERE clause that behaves the same way as in the SELECT expression. The result of the DELETE statement depends on whether a RETURNING clause is present or not. Without a RETURNING clause the DELETE returns the number of rows deleted. Otherwise, for each deleted row the expressions following the RETURNING clause are computed the same way as in the SELECT clause and the result is returned to the application. Finally, the DELETE statement may start with declarations of external variables used in the rest of the statement. As in queries, such declarations are mandatory.

If any error occurs during the execution of a DELETE statement, there is a possibility that some rows will be deleted and some not. The system does not keep track of what rows got deleted and what rows are not yet deleted. This is because Oracle NoSQL Database focuses on low latency operations. Long running operations across shards are not coordinating using two-phase commit and lock mechanism. In such cases, it is recommended that the application re-run the DELETE statement.

Example 7-24 Deleting Rows with SQL

The following statement deletes all users whose age is less than 16, returning the first and last name of each deleted user.

```
DELETE FROM users
WHERE age < 16
RETURNING firstName, lastName
```

Example 7-25 Delete from the storeAcctComposite table the shopper's data related to fulfilled orders in the past year

You can use the DELETE statement to remove data from a JSON collection table. The DELETE operation works similarly on fixed schema tables. Consider the JSON collection table created for a [shopping application](#).

```
DELETE FROM storeAcctComposite s where s.orders.EstDelivery <any "2023-01-01"
AND s.orders.status =any "Delivered" RETURNING contactPhone, id, firstName
```

Explanation: In the storeAcctComposite table, the shoppers don't have a permanent account. The contactPhone and id are used to track the orders. As a maintenance activity, the DELETE statement above deletes the shoppers' data for which all the orders are already delivered and the delivery dates are before the year 2023.

Output:

```
{
  "contactPhone" : "1517113582",
  "id" : 10,
  "firstName" : "Dierdre"
}
```

UPDATE Statement

You can use an update statement to update a row in a table. It also supports updating multiple rows in a table that share the same shard key. You can find which columns in your table comprise of the shard key [here](#), if you are running an On-Premises version of Oracle NoSQL.

Syntax

```
update_statement ::=
  UPDATE table\_name [[AS] table\_alias]
  update\_clause ("," update_clause)*
  WHERE expression
  [returning_clause]

returning_clause ::= RETURNING select\_list
```

Semantics

The update takes place at the server, eliminating the read-modify-write cycle, that is, the need to fetch the whole row at the client, compute new values for the targeted fields (potentially based on their current values) and then send the whole row back to the server.

Both syntactically and semantically, the update statement of the Oracle NoSQL Database is similar to the update statement of standard SQL, but with extensions to handle the richer data model of the Oracle NoSQL Database. So, as shown by the syntax above:

- First, you specify the table to update by its name and an optional table alias (the alias may be omitted only if top-level columns only are to be accessed; otherwise, as in read-only queries, the alias is required as the first step of path expressions that access nested fields).
- Then come one or more update clauses.

- The WHERE clause specifies what rows to update. The current implementation supports both single-row and multiple-row updates
 - To update a single row, the WHERE clause must specify the complete primary key.
 - To update multiple rows, the WHERE clause must specify the shard key, ensuring that all rows reside on the same shard. This enables Oracle NoSQL Database to perform this update in an ACID transaction. For examples, see [Example: Updating multiple rows](#).

A single update transaction applies all changes if successful or makes no change if an error occurs.
- There is an optional RETURNING clause. The RETURNING clause works only when the WHERE clause includes the complete primary key. Currently, it only supports single-row updates.
 - If the clause is not present, the result of the update statement is the number of rows updated. For example,
 - * Zero is returned if there is no row satisfying the conditions in WHERE clause, or if the updates specified by the update clauses turned out to be no-ops for the row selected by the WHERE clause.
 - * One is returned if there is a single row satisfying the conditions in the WHERE clause, and so on.
 - Otherwise, if there is a RETURNING clause, it acts the same way as the SELECT clause: it can be a "*", in which case, the full updated row will be returned, or it can have a list of expressions specifying what needs to be returned.
- There is also a limit on the number of records a single update query can update. By default, the system sets this limit to 1,000. If an update statement tries to modify more than 1,000 records in a single transaction, it will fail during the process. To avoid this, use `setLimit(int limit)` method, to override the default limit value. For more information, see `QueryRequest.setLimit(int limit)`. By setting the desired limit, you can work with large update requirements.

Note

While you can override the default limit, we do not recommend doing so for a large number of rows, as this can cause the update operation or the other read and write operations that affect the same data during the transaction to time out or experience higher latency.

The UPDATE statement is used to update the rows in the table. If you want to supply user-supplied metadata for the update operation, you can use the available `QueryRequest` API. For an example, see [Using User-Supplied Metadata in Write Operations](#) from *Developers Guide*.

Update Clauses

Syntax

```
update_clause ::=
  (SET set\_clause ("," (update_clause | set_clause))* ) |
  (ADD add\_clause ("," (update_clause | add_clause))* ) |
  (PUT put\_clause ("," (update_clause | put_clause))* ) |
  (REMOVE remove\_clause ("," remove_clause)* ) |
  (JSON MERGE JSON MERGE Clause ("," (update_clause |
```

```
json_merge_patch_clause))* |  
(SET TTL ttl\_clause ("," update_clause)*)
```

Semantics

There are six kinds of update clauses:

SET

Updates the value of one or more existing fields. See [SET Clause](#).

ADD

Adds new elements to one or more arrays. See [ADD Clause](#).

PUT

Adds new fields to one or more maps. It can also update the values of existing map fields. See [PUT Clause](#).

REMOVE

Removes elements/fields from one or more arrays/maps. See [REMOVE Clause](#).

JSON MERGE

Updates JSON documents with the changes specified in the patch. See [JSON MERGE Clause](#)

SET TTL

Updates the expiration time of the row. See [SET TTL Clause](#).

The update clauses are applied immediately, in the order they appear in the update statement, so the effects of each clause are visible to subsequent clauses. Although the syntax allows for multiple SET TTL clauses, only the last one will be effective; the earlier ones, if any, are ignored.

The SET, ADD, PUT, and REMOVE clauses start with a target expression, which computes the items to be updated or removed. In all cases, the target expression must be either a top-level column reference or a path expression starting with the table alias. If the target expression returns nothing, the update clause is a no-op.

SET Clause

Syntax

```
set_clause ::= path\_expression "=" expression
```

Semantics

The SET clause consists of two expressions: the target expression and the new-value expression. The target expression returns the items to be updated. Notice that a target item may be atomic or complex, and it will always be nested inside a complex item (its parent item). For each such target item, the new-value expression is evaluated, and its result replaces the target item within the parent item.

If the target expression returns a NULL item, then either the target item itself is the NULL item, or one of its ancestors is NULL. In the former case, the target item will be replaced by the new item. In the latter case the SET is a no-op.

The new-value expression may return zero or more items. If it returns an empty result, the SET is a no-op. If it returns more than one item, the items are enclosed inside a newly constructed

array (this is the same as the way the SELECT clause treats multi-valued expressions in the select list). So, effectively, the result of the new-value expression contains at most one item. This new item is then cast to the type expected by the parent item for the target field. This cast behaves like the cast expression as described in the [Cast Expression](#) section. If the cast fails, an error is raised; otherwise, the new item replaces the target item within the parent item.

The new-value expression may reference the implicitly declared variable \$, which is bound to the current target item. Use of the \$ variable makes it possible to have target expressions that return more than one item. As mentioned already, in this case, the SET clause will iterate over the target items, and for each target item T, bind the \$ variable to T, compute the new-value expression, and replace T with the result of the new-value expression.

What if the new-value expression is the (reserved) keyword null? Normally, null is interpreted as the json null value. However, if the parent of the target item is a record, then null will be interpreted as the SQL NULL, and the targeted record field will be set to the SQL NULL.

See [Example: Updating Rows](#).

ADD Clause

Syntax

```
add_clause ::=  
    path\_expression [add\_expression] expression
```

Semantics

The ADD clause is used to add new elements into one or more arrays. It consists of a target expression, which should normally return one or more array items, an optional position expression, which specifies the position within each array where the new elements should be placed, and a new-elements expression that returns the new elements to insert.

The ADD clause iterates over the sequence returned by the target expression. For each target item, if the item is not an array it is skipped. Otherwise, the position expression (if present) and the new-elements expression are computed for the current target array. These two expressions may reference the \$ variable, which is bound to the current target array.

If the new-values expression returns nothing, the ADD is a no-op. Otherwise, each item returned by this expression is cast to the element type of the array. An error is raised if any of these casts fail. Otherwise, the new elements are inserted into the target array, as described below.

If the position expression is missing, or if it returns an empty result, the new elements are appended at the end of the target array. An error is raised if the position expression returns more than one item or a non-numeric item. Otherwise, the returned item is cast to an integer. If this integer is less than 0, it is set to 0. If it is greater or equal to the array size, the new elements are appended. Otherwise, if the integer position is P and the new-elements expression returns N items, the 1st item is inserted at position P, the 2nd at position P+1, and so on. The existing array elements at position P and afterwards are shifted N positions to the right.

See [Example: Updating Rows](#).

PUT Clause

Syntax

`put_clause ::= path\_expression expression`

Semantics

The PUT clause is used primarily to add new fields into one or more maps. It consists of a target expression, which should normally return one or more map item and a new-fields expression that returns one or more maps or records, whose fields are inserted in the target maps.

The PUT clause iterates over the sequence returned by the target expression. For each target item, if the item is not a map it is skipped. Otherwise, the new-fields expression is computed for the current target map. The new-maps expression may reference the \$ variable, which is bound to the current target map.

If the new-fields expression returns nothing, the PUT is a no-op. Otherwise, for each item returned by the new-fields expression, if the item is not a map or a record, it is skipped, else, the fields of the map/record are "merged" into the current target map. This merge operation will insert a new field into the target map if the target map does not already have a field with the same key; Otherwise, it will set the value of the target field to the value of the new field.

See [Example: Updating Rows](#).

REMOVE Clause

Syntax

`remove_clause ::= path\_expression`

Semantics

The remove clause consists of a single target expression, which computes the items to be removed. The REMOVE clause iterates over the target items. For each such item, if its parent is a record, an error is raised. Otherwise, if the target item is not NULL, it is removed from its parent. If the target item is NULL, then since arrays and map cannot contain NULLs, one of its ancestors must be NULL. In this case, the NULL is skipped.

See [Example: Updating Rows](#).

JSON MERGE Clause

Syntax

`json_merge_patch_clause ::= target_expr WITH PATCH json_patch_expr
json_patch_expr ::= map_constructor | array_constructor | constant_expression |
variable_reference`

Semantics

The JSON MERGE clause conforms to the standard specified in RFC 7396 . This clause allows you to make a set of changes to a JSON document in a single statement. The JSON MERGE clause consists of two expressions: the target expression and the patch expression.

The target expression specifies the JSON document that needs update. When evaluated on a specified table row, this expression must return zero or more [valid JSON values](#). Otherwise, the JSON MERGE clause is a no-op.

The patch expression, when evaluated, must return a single valid JSON value. This document describes the updates to be made to each of the JSON values returned by the target JSON document. If the evaluated patch contains a JSON value that does not exist within the target, the path and its associated value are added. If the target contains the path, its value is replaced with the value in the patch. If the patch contains a path with a null value and the path exists in the target, the field in the path is removed from the target. The updates are made according to the algorithm described in RFC 7396. In summary, if the patch expression returns an empty result, the JSON MERGE clause is a no-op.

The patch expression can be a constant literal, reference to an external variable, map, or an array constructor. The patch expression must not access any table data as the SQL engine evaluates the patch expression before the execution of the UPDATE statement.

You can use the JSON MERGE patch with all other update clauses (SET, ADD, PUT, REMOVE) in a single UPDATE statement.

Note

- If the patch expression includes values other than a JSON object, the result replaces the entire target with the entire patch value.
- If the patch expression includes a JSON object, the entire JSON object is either created (if it does not exist) or replaced in the target.
- If the patch expression includes a JSON path that does not exist in the target, the path is created.
- The target can include path expression of either parent-level or nested-level JSON object.

See [Example: Updating JSON Data](#).

SET TTL Clause

Syntax

```
ttn_clause ::=  
  (add expression (HOURS | DAYS)) |  
  (USING TABLE DEFAULT)
```

Semantics

If a SET TTL clause is used with an UPDATE statement, a new expiration time is computed and applied to the row being updated. In case of MR Tables, the rows replicated to other regions carry the recalculated expiration time of the row being updated. Therefore, this row will

have the same expiration time in all the regions after successful replication. If a TTL value is updated to the same row in multiple regions, then the TTL value will be set in all regions to the value held in the row with the greatest write timestamp.

See [Example: Updating TTL](#).

Updating rows with an IDENTITY Column

An IDENTITY column that is defined as GENERATED ALWAYS cannot be updated. Only IDENTITY column that is defined as GENERATED BY DEFAULT can be updated.

Updating rows with an MR_COUNTER column

You can update an MR_COUNTER column in a multi-region table by incrementing or decrementing the values using standard arithmetic computations. For creating a table with MR_COUNTER column, see [Create table using MR_COUNTER datatype](#)

Note

The system will return an error if you use the UPDATE clauses on MR_COUNTERS in the following scenarios:

- SET or PUT clauses to explicitly supply a value to an MR_COUNTER in the table.
- REMOVE clause to remove an MR_COUNTER column from the table.

You can use an ALTER statement to drop an MR_COUNTER column from the table. For more details, see [Add or Remove an MR_COUNTER column](#).

Incrementing the value of an MR_COUNTER column.

A simple example would be incrementing the likes a user gets on a social media website.

Example 7-26 Increment an MR_COUNTER value

```
CREATE Table myTable( name STRING,  
                      count INTEGER AS MR_COUNTER,  
                      PRIMARY KEY(name)) IN REGIONS DEN,LON
```

```
INSERT INTO myTable(name) VALUES ('Chris')
```

```
UPDATE myTable SET count = count + 10 WHERE name = "Chris"
```

Decrementing the value of an MR_COUNTER column.

The following example decrements the value of an MR_COUNTER.

Example 7-27 Decrement an MR_COUNTER value

```
CREATE Table myTable( name STRING,
                      count INTEGER AS MR_COUNTER,
                      PRIMARY KEY(name)) IN REGIONS DEN,LON
```

```
INSERT INTO myTable VALUES ('Chris',10)
```

```
UPDATE myTable SET count = count - 4 WHERE name = "Chris"
```

Updating JSON MR_COUNTER values

You can update a JSON MR_COUNTER column (the same way as an MR_Counter column) in a multi-region table.

Example 7-28 Increment the value of a JSON MR_COUNTER column

You can update a JSON MR_COUNTER column by incrementing the value.

```
UPDATE demoJSONMR a SET a.jsonWithCounter.counter = a.jsonWithCounter.counter
+ 1
WHERE name = "Anna"
```

You can also update a JSON MR_COUNTER column by decrementing the value.

```
UPDATE demoJSONMR a SET a.jsonWithCounter.counter = a.jsonWithCounter.counter
- 1
WHERE name = "Anna"
```

Updating counter values in both regions and perform a merge

When MR_COUNTER fields exist in both the remote JSON field and local JSON field, the system merges them as MR_COUNTER data types. You can update the MR_COUNTER fields in the remote and local region independently. The system automatically performs a merge on these concurrent modifications without user intervention.

Example 7-29 Update MR_COUNTER fields in the remote and local region

Consider the table demoJSONMR created in regions FRA and LON with the same definition as shown below.

```
CREATE TABLE demoJSONMR(name STRING,
                          jsonWithCounter JSON(counter as INTEGER MR_COUNTER,
                                                  person.count as LONG MR_COUNTER),
                          PRIMARY KEY(name)) IN REGIONS FRA,LON
```

Step 1: Insert a row into the demoJSONMR table in the region FRA.

```
INSERT INTO demoJSONMR VALUES (
    Anna,
    {
        "id" : 1,
        "counter" : NULL,
```

```

        "person" : {
            "age" : 10,
            "number" : 100
        }
    }
}

```

Step 2: Update the row inserted above and increment JSON MR_COUNTER field "counter".

```

UPDATE demoJSONMR a SET a.jsonWithCounter.counter = a.jsonWithCounter.counter
+ 1
WHERE name = "Anna"

```

Step 3: In the Remote region LON, insert a row into the table with the same primary key "Anna", but different values for other fields.

```

INSERT INTO exampleTable VALUES (
    Anna,
    {
        "id" : 2,
        "counter" : NULL,
        "person" : {
            "age" : 10,
            "number" : 101
        }
    }
}

```

Step 4: In the Remote region LON, update the record and increment the JSON MR_COUNTER field "counter".

```

UPDATE demoJSONMR a SET a.jsonWithCounter.counter = a.jsonWithCounter.counter
+ 1
WHERE name = "Anna"

```

Step 5: In the statement above, the remote row gets updated. This update gets merged with the local row and the field "counter", as shown below.

```

SELECT * FROM demoJSONMR WHERE name = "Anna"

```

Output:

```

{"name":"Anna",
"jsonWithCounter":{"counter":2,"id":2,"person":
    {"age":10,"count":0,"number":101}}
}

```

If the remote JSON and local JSON for MR_COUNTER have mismatched definitions, the INSERT or UPDATE operation is not performed on the mismatched schema. These rows are logged as incompatible rows.

Updating MR_COUNTER values in JSON collection table

Consider a JSON collection table created for a [shopping application with MR_COUNTER](#). The mycounter field in the table is an MR_COUNTER with its value set to zero upon inserting data into the table.

Example 7-30 Update the MR_COUNTER value in a JSON collection table

The following is a sample row from the shopping application table:

```
{ "contactPhone": "1817113382", "address":
{ "city": "Houston", "number": 401, "state": "TX", "street": "Tex
Ave", "zip": 95085 }, "cart": [ { "item": "handbag", "priceperunit": 350, "quantity": 1 },
{ "item": "Lego", "priceperunit": 5500, "quantity": 1 } ], "firstName": "Adam", "lastName":
"Smith", "mycounter": 0 }
```

To update an MR_COUNTER value, you must supply the MR_COUNTER's name in the SET clause as illustrated in the query below:

```
UPDATE storeAcctMR s SET s.mycounter = s.mycounter + 5 WHERE
s.contactPhone="1817113382"
```

In this example, you increment the value of the MR_COUNTER by 5 for the shopper with the contact number "1817113382". You get the following output when you fetch the shopper's record:

```
{ "contactPhone": "1817113382", "address":
{ "city": "Houston", "number": 401, "state": "TX", "street": "Tex
Ave", "zip": 95085 }, "cart": [ { "item": "handbag", "priceperunit": 350, "quantity": 1 },
{ "item": "Lego", "priceperunit": 5500, "quantity": 1 } ], "firstName": "Adam", "lastName":
"Smith", "mycounter": 5 }
```

Example: Updating Rows

Let's assume a table, called People, with only two columns: an integer id column and an info column of type JSON. Furthermore, let's assume the following row to be updated:

```
CREATE TABLE People (
    id INTEGER,
    info JSON,
    PRIMARY KEY(id))
```

```
INSERT INTO People VALUES (
    0,
    {
        "firstName": "John",
        "lastName": "Doe",
        "profession": "software engineer",
        "income": 200000,
        "address": {
            "city": "San Francisco",
            "state": "CA",
```

```

        "phones" : [
            { "areacode":415, "number":2840060, "kind":"office" },
            { "areacode":650, "number":3789021, "kind":"mobile" },
            { "areacode":415, "number":6096010, "kind":"home" }
        ]
    },
    "children": {
        "Anna" : {
            "age" : 10,
            "school" : "school_1",
            "friends" : ["Anna", "John", "Maria"]
        },
        "Ron" : { "age" : 2 },
        "Mary" : {
            "age" : 7,
            "school" : "school_3",
            "friends" : ["Anna", "Mark"]
        }
    }
}
)

```

Example 7-31 Use UPDATE statement to update various fields in the table

The following update statement updates various fields in the above row:

```

UPDATE People p
  SET p.info.profession = "surfing instructor",
  SET p.info.address.city = "Santa Cruz",
  SET p.info.income = p.info.income / 10,
  SET p.info.children.values().age = $ + 1,
  ADD p.info.address.phones
    0 { "areacode":831, "number":5294368, "kind":"mobile" },
  REMOVE p.info.address.phones [$element.kind = "office"],
  PUT p.info.children.Ron { "friends" : ["Julie"] },
  ADD p.info.children.values().friends seq_concat("Ada", "Aris")
WHERE id = 0
RETURNING *

```

After the update, the row looks like this:

```

{
  "id":0,
  "info":{
    "firstName":"John",
    "lastName":"Doe",
    "profession":"surfing instructor",
    "income":20000,
    "address":{
      "city":"Santa Cruz",
      "phones":[
        {"areacode":831,"kind":"mobile","number":5294368},
        {"areacode":650,"kind":"mobile","number":3789021},
        {"areacode":415,"kind":"home","number":6096010}
      ]
    }
  }
}

```

```

        "state": "CA"
    },
    "children": {
        "Anna": {
            "age": 11,
            "friends": ["Anna", "John", "Maria", "Ada", "Aris"],
            "school": "school_1"
        },
        "Ron": {
            "age": 3,
            "friends": ["Julie", "Ada", "Aris"]
        },
        "Mary": {
            "age": 8,
            "friends": ["Anna", "Mark", "Ada", "Aris"],
            "school": "school_3"
        }
    }
}

```

The first two SET clauses change the profession and city of John Doe. The third SET reduces his income to one-tenth. The fourth SET increases the age of his children by 1. Notice the use of the \$ variable here: the expression `p.info.children.values().age` returns 3 ages; The SET will iterate over these ages, bind the \$ variable to each age in turn, compute the expression `$ + 1` for each age, and update the age with the new value. Notice that the income update can also have used a \$ variable: `set p.info.income = $ / 10`. This would have saved the re-evaluation of the `p.info.income` path on the right-hand side or the `"="`.

The ADD clause adds a new phone at position 0 inside the phones array. The REMOVE removes all the office phones (only one in this example). The PUT clause adds a friend for Ron. In this clause, the expression `p.info.children.Ron` returns the value associated with the Ron child. This value is a map (the json object `{ "age" : 3 }`) and becomes the target of the update. The 2nd expression in the PUT (`{ "friends" : ["Julie"] }`) constructs and returns a new map. The fields of this map are added to the target map. Finally, the last ADD clause adds the same two new friends to each child. See [seq_concat function](#) function.

Notice that the update query in this example would have been exactly the same if instead of type JSON, the info column had the following RECORD type:

```

RECORD(
    firstName STRING,
    lastName STRING,
    profession STRING,
    income INTEGER,
    address RECORD(
        city STRING,
        state STRING,
        phones ARRAY(
            RECORD(
                areacode INTEGER,
                number INTEGER,
                kind STRING
            )
        )
    ),
)

```

```

        children MAP(
            RECORD(
                age INTEGER,
                school STRING,
                friends ARRAY(STRING)
            )
        )
    )
)

```

Example: Updating multiple rows

Create a table named `EmployeeInfo` with an integer `empID` column as the primary key, a string `department` column as the shard key, a string `fullName` column, and a JSON type `info` column.

```

CREATE TABLE EmployeeInfo (
    empID INTEGER,
    department STRING,
    fullName STRING,
    info JSON,
    PRIMARY KEY(SHARD(department), empID))

```

Next, add the following rows:

```

INSERT INTO EmployeeInfo VALUES (101, "HR", "Liam Phillips",
{"salary":100000,"address":{"city":"Toronto","country":"Canada"}})

INSERT INTO EmployeeInfo VALUES (102, "HR", "Emma Johnson",
{"salary":105000,"address":{"city":"Melbourne","country":"Australia"}})

INSERT INTO EmployeeInfo VALUES (103, "IT", "Carlos Martinez",
{"salary":110000,"address":{"city":"Barcelona","country":"Spain"}})

INSERT INTO EmployeeInfo VALUES (104, "Finance", "Sophia Becker",
{"salary":130000,"address":{"city":"Munich","country":"Germany"}})

INSERT INTO EmployeeInfo VALUES (105, "IT", "Ethan Dalmini",
{"salary":250000,"address":{"city":"Cape Town","country":"South Africa"}})

INSERT INTO EmployeeInfo VALUES (106, "HR", "James Peterson",
{"salary":100500,"address":{"city":"Istanbul","country":"Turkey"}})

INSERT INTO EmployeeInfo VALUES (107, "Finance", "John Doe",
{"salary":250000,"address":{"city":"Seoul","country":"South Korea"}})

```

Example 7-32 Use UPDATE statement to update multiple rows in the table

The following statement updates the specified field in the rows associated with the mentioned shard key.

```

UPDATE EmployeeInfo emp
SET emp.info.address.city="Oslo",
SET emp.info.address.country="Norway",

```

```
SET emp.info.salary = emp.info.salary + 5000
where department="HR"
```

Explanation: In the above query, the SET clause updates the `emp.info.address.city` field to 'Oslo,' the `emp.info.address.country` field to 'Norway,' and increases the `emp.info.salary` field by 5000 for all rows where the department column, designated as the shard key, equals 'HR.' Because the UPDATE statement only mentions the shard key, the database returns only the number of rows it updates.

Output:

```
+-----+
| NumRowsUpdated |
+-----+
|          3      |
+-----+
```

Now, run the SELECT query to verify the updated rows. Check that you have updated the `emp.info.address.city` and `emp.info.address.country` fields to 'Oslo' and 'Norway,' respectively, and confirm that you have increased the `emp.info.salary` field by 5000 for all employees working in the HR department.

```
select * from EmployeeInfo
```

Output:

empID	department	fullName	info
103	IT	Carlos Martinez	address city Barcelona country Spain salary 110000
105	IT	Ethan Dalmini	address city Cape Town country South Africa salary 250000
101	HR	Liam Phillips	address city Oslo country Norway salary 105000
102	HR	Emma Johnson	address city Oslo country Norway salary 110000
106	HR	James Peterson	address city Oslo country Norway salary 105500

104	Finance	Sophia Becker	address	city	Munich
			country	Germany	
			salary	130000	
107	Finance	John Doe	address	city	Seoul
			country	South Korea	
			salary	250000	

7 rows returned

Example: Updating JSON Data

You can update JSON data using either JSON MERGE clause or ADD clause.

JSON MERGE Clause

You can use the [JSON MERGE clause](#) to make a set of changes to a JSON document. The samples in the following section consider **airline baggage tracking application** and **streaming media service application** from the [CREATE TABLE](#) topic.

Example 7-33 Update existing fields and add a new field to a JSON document

```
UPDATE BaggageInfo b
JSON MERGE b.bagInfo[0] WITH PATCH {"tagNum" : "18657806255240", "routing" :
"MIA/LAX", "lastSeenStation" : "LAX", "bagStatus" : "delivered"}
WHERE ticketNo = 1762344493810 RETURNING b.bagInfo[0].tagNum,
b.bagInfo[0].routing, b.bagInfo[0].lastSeenStation, b.bagInfo[0].bagStatus
```

Explanation: In the query above, you update a passenger's record in the airline baggage tracking application to correct a few inadvertent errors. The JSON MERGE clause compares the content of the patch with the current contents and updates the tagNum, routing, and lastSeenStation fields in the first element of the bagInfo JSON array. As the bagStatus field is currently not a part of the first element in the bagInfo array, the JSON MERGE patch inserts it into the document.

You use the UPDATE statement with a [RETURNING clause](#) to display only the updated fields.

Output:

```
{
  "tagNum" : "18657806255240",
  "routing" : "MIA/LAX",
  "lastSeenStation" : "LAX",
  "bagStatus" : "delivered"
}
```


Note

Here, bagInfo is an array of JSON documents, which represents the checked bags per passenger. If you don't specify the array element to be updated, the JSON MERGE clause replaces the entire bagInfo array with the patch content. If there are more than one checked bag per passenger, you can use the JSON MERGE clause iteratively in the same UPDATE statement as follows:

```
UPDATE BaggageInfo b
JSON MERGE b.bagInfo[0] WITH PATCH {"tagNum" : "18657806255111",
"routing" : "MIA/LAX", "lastSeenStation" : "LAX", "bagStatus" :
"delivered"},
JSON MERGE b.bagInfo[1] WITH PATCH {"tagNum" : "18657806255112",
"routing" : "MIA/LAX", "lastSeenStation" : "LAX", "bagStatus" :
"delivered"}
WHERE ticketNo = 1762320369957
```

Example 7-34 Add a new field using its JSON path in the patch

Consider the following passenger's data from the airline baggage tracking application. The nested flightLegs document does not include the flight number.

```
{"ticketNo":1882344493810,"fullName":"Joan
Smith","gender":"F","contactPhone":"886-324-1064","confNo":"LE6F4Z","bagInfo":
[{"flightLegs":
{"estimatedArrival":"2019-02-01T11:00:00Z","flightDate":"2019-02-01T06:00:00Z"
,"fltRouteDest":"LAX","fltRouteSrc":"MIA"},"id":"79039899165297","routing":"MI
A/LAX","tagNum":"17657806255240"}]}
```

You can apply a patch to include the flightNo field into the flightLegs document as follows:

```
UPDATE BaggageInfo b
JSON MERGE b.bagInfo[0] WITH PATCH {"flightLegs": {"flightNo":"BM107"}} WHERE
ticketNo = 1882344493810 RETURNING *
```

Explanation: Here, you supply the JSON path of the flightNo field in the patch expression. The JSON MERGE clause evaluates the JSON path and inserts the flightNo field into the nested flightLegs document of the first element of the bagInfo array.

Output:

```
{"ticketNo":1882344493810,"fullName":"Joan
Smith","gender":"F","contactPhone":"886-324-1064","confNo":"LE6F4Z","bagInfo":
[{"flightLegs":
{"estimatedArrival":"2019-02-01T11:00:00Z","flightDate":"2019-02-01T06:00:00Z"
,"flightNo":"BM107","fltRouteDest":"LAX","fltRouteSrc":"MIA"},"id":"7903989916
5297","routing":"MIA/LAX","tagNum":"17657806255240"}]}
```

Note

If the nested flightLegs document does not exist, the JSON MERGE patch creates one.

Example 7-35 Remove field from a JSON document and insert a nested document

```
UPDATE stream_acct s
JSON MERGE s.acct_data WITH PATCH {"country" : NULL, "recommended" :
{"showName1" : "laugh it", "showName2": "Mystery solved", "showName3":
"bakesnCooks"}}
WHERE acct_id = 1 RETURNING *
```

Explanation: In the query above, you use the JSON MERGE clause to update a subscriber record in the streaming media service application. In the patch expression, you supply the country and recommended fields. The JSON MERGE patch inserts the new nested recommended JSON document into the acct_data JSON document. You indicate the removal of the country field from the document by supplying a NULL value.

The RETURNING clause fetches the entire updated row.

Output:

```
{
  "acct_id" : 1,
  "profile_name" : "AP",
  "account_expiry" : "2023-10-18T00:00:00.000000000Z",
  "acct_data" : {
    "contentStreamed" : [{
      "genres" : ["action", "crime", "spanish"],
      "numSeasons" : 4,
      "seriesInfo" : [{
        "episodes" : [{ .... }],
        "showId" : 26,
        "showName" : "At the Ranch",
        "showtype" : "tvseries"
      }], {
        "genres" : ["comedy", "french"],
        "numSeasons" : 2,
        "seriesInfo" : [{
          "episodes" : [{ .... }],
          "showId" : 15,
          "showName" : "Bienvenu",
          "showtype" : "tvseries"
        }],
        "firstName" : "Adam",
        "lastName" : "Phillips",
        "recommended" : {
          "showName1" : "laugh it",
          "showName2" : "Mystery solved",
          "showName3" : "bakesnCooks"
        }
      }
    ]
  }
}
```

```
    }
  }
```

Example 7-36 Update/add a field to a document

Consider following subscriber records from the streaming media service application.

```
{ "acct_id":1, "profile_name": "AP", "account_expiry": "2023-10-18T00:00:00.000000000Z", "acct_data": { "contentStreamed": [ { "genres": [ "action", "crime", "spanish", "numSeasons":4, "showId":26, "showName": "At the Ranch", "showtype": "tvseries" }, { "genres": [ "comedy", "french", "numSeasons":2, "showId":15, "showName": "Bienvenu", "showtype": "tvseries" ] }, "firstName": "Adam", "lastName": "Phillips" } }
```

```
{ "acct_id":2, "profile_name": "AW", "account_expiry": "2025-12-18T00:00:00.000000000Z", "acct_data": { "contentStreamed": { "genres": [ "comedy", "french", "numSeasons":2, "showId":15, "showName": "Bienvenu", "showtype": "tvseries" ], "country": "France", "firstName": "Adelaide", "lastName": "Willard" } }
```

Example 4a: Use the JSON MERGE clause to update the country field in the acct_data JSON document for the records above:

```
UPDATE stream_acct s
JSON MERGE s.acct_data WITH PATCH { "country" : "USA" } WHERE acct_id = 1
RETURNING s.acct_id, s.acct_data.country
```

```
UPDATE stream_acct s
JSON MERGE s.acct_data WITH PATCH { "country" : "Italy" } WHERE acct_id = 2
RETURNING s.acct_id, s.acct_data.country
```

Explanation: In the first UPDATE statement for the acct_id 1, the JSON MERGE clause inserts the country field to the acct_data JSON document as the field does not exist in the original document. In the second UPDATE statement, the JSON MERGE clause updates the country field value from France to Italy for the subscriber acct_id 2.

Output:

```
{ "acct_id":1, "country": "USA" }
1 row returned
```

```
{ "acct_id":2, "country": "Italy" }
1 row returned
```

Example 4b: Update the country field using path expression.

Consider the original subscriber records from the streaming media service application.

Update the country field using its path expression for the record with the acct_id 2.

```
UPDATE stream_acct s
JSON MERGE s.acct_data.country WITH PATCH "Switzerland" where acct_id = 2

{"NumRowsUpdated":1}
```

Explanation: The UPDATE operation is successful, and the country field value is replaced with the supplied value. Now update the subscriber record for acct_id 1.

```
UPDATE stream_acct s
JSON MERGE s.acct_data.country WITH PATCH "USA" where acct_id = 1

{"NumRowsUpdated":0}
```

In contrast to the UPDATE operation in the Example 4a for acct_id 1, where the country field was successfully inserted, here it is a no-op. When the field that is directly accessed using its path expression does not exist in the target JSON document, the resulting JSON MERGE operation is a no-op.

Example 7-37 Update a JSON NULL field

Consider the following subscriber record from the streaming media service application.

```
{"acct_id":1,"profile_name":"AP","account_expiry":"2023-10-18T00:00:00.000000000Z","acct_data":null}
```

Here, the acct_data is a JSON field with a NULL value. Use the JSON MERGE clause to update the field.

```
JSON MERGE s.acct_data WITH PATCH {"contentStreamed":{"genres":
["comedy","french"],"numSeasons":2,"showId":15,"showName":"Bienvenu","showtype":
"tvseries"},"country":NULL,"firstName":"Adam","lastName":"null"} WHERE
acct_id = 1 RETURNING *
```

Explanation: The JSON MERGE patch replaces the NULL value in the acct_data document with the supplied JSON values from the patch expression. Observe that you supply a JSON NULL value to the country field. The JSON MERGE patch does not insert this field in the acct_data document. However, the lastName field with a string value null is successfully inserted.

Output:

```
{"acct_id":1,"profile_name":"AP","account_expiry":"2023-10-18T00:00:00.000000000Z","acct_data":{"contentStreamed":{"genres":
["comedy","french"],"numSeasons":2,"showId":15,"showName":"Bienvenu","showtype":
"tvseries"},"firstName":"Adam","lastName":"null"}}
```

ADD clause

You can use [ADD clause](#) to update a field in a JSON document. You can add one or more array elements to an existing array using the ADD clause. You can also optionally indicate the position of the new elements to be added in the array.

Example 7-38 Adding a single element to an existing array in JSON data

The People table has one row currently as shown below. For creating the table and inserting data, see [Example: Updating Rows](#).

```
SELECT * FROM People
```

```
{
  "id" : 0,
  "info" : {
    "address" : {
      "city" : "Santa Cruz",
      "phones" : [{
        "areacode" : 831,
        "kind" : "mobile",
        "number" : 5294368
      }, {
        "areacode" : 650,
        "kind" : "mobile",
        "number" : 3789021
      }, {
        "areacode" : 415,
        "kind" : "home",
        "number" : 6096010
      }
    ],
    "state" : "CA"
  },
  "children" : {
    "Anna" : {
      "age" : 11,
      "friends" : ["Anna", "John", "Maria", "Ada", "Aris"],
      "school" : "school_1"
    },
    "Mary" : {
      "age" : 8,
      "friends" : ["Anna", "Mark", "Ada", "Aris"],
      "school" : "school_3"
    },
    "Ron" : {
      "age" : 3,
      "friends" : ["Julie", "Ada", "Aris"]
    }
  },
  "firstName" : "John",
  "income" : 20000,
  "lastName" : "Doe",
  "profession" : "surfing instructor"
}
```

Add a new element to the phones array at the beginning of the array.

```
UPDATE People p ADD p.info.address.phones 0
{"areacode":499, "number":33864368, "kind":"mobile" }
WHERE id = 0
```

```
SELECT * FROM People
```

Output:

```
{
  "id" : 0,
  "info" : {
    "address" : {
      "city" : "Santa Cruz",
      "phones" : [{
        "areacode" : 499,
        "kind" : "mobile",
        "number" : 33864368
      }, {
        "areacode" : 831,
        "kind" : "mobile",
        "number" : 5294368
      }, {
        "areacode" : 650,
        "kind" : "mobile",
        "number" : 3789021
      }, {
        "areacode" : 415,
        "kind" : "home",
        "number" : 6096010
      }
    ],
    "state" : "CA"
  },
  "children" : {
    "Anna" : {
      "age" : 11,
      "friends" : ["Anna", "John", "Maria", "Ada", "Aris"],
      "school" : "school_1"
    },
    "Mary" : {
      "age" : 8,
      "friends" : ["Anna", "Mark", "Ada", "Aris"],
      "school" : "school_3"
    },
    "Ron" : {
      "age" : 3,
      "friends" : ["Julie", "Ada", "Aris"]
    }
  },
  "firstName" : "John",
  "income" : 20000,
```

```
        "lastName" : "Doe",  
        "profession" : "surfing instructor"  
    }  
}
```

Example 7-39 Adding an array of elements to an existing array in JSON data

When you need to add more than one element of an array to an existing array in JSON data, you need to add the new-elements expressions inside parentheses and optionally add the position expression (if any).

The following query throws an error as shown below:

```
UPDATE People p  
ADD p.info.address.phones  
0 { "areacode":5, "number":1, "kind":"mobile" },  
  { "areacode":6, "number":2, "kind":"mobile" }  
WHERE id = 0
```

Output:

```
Error handling command UPDATE People p  
ADD p.info.address.phones  
0 { "areacode":5, "number":1, "kind":"mobile" },  
  { "areacode":6, "number":2, "kind":"mobile" }  
WHERE id = 0:  
Error: at (5, 12) mismatched input '<EOF>' expecting {WHERE, ',', '}', at line  
5:12  
rule stack: [parse, statement, update_statement]
```

This can be corrected using the one of the two different options as shown below:

Option 1:

```
UPDATE People p  
ADD p.info.address.phones  
([ { "areacode":1, "number":1, "kind":"mobile" },  
  { "areacode":2, "number":2, "kind":"mobile" }  
  ] )  
WHERE id = 0
```

Output:

```
{  
  "NumRowsUpdated" : 1  
}  
1 row returned
```

Option 2:

```
UPDATE People p  
ADD p.info.address.phones
```

```
0 [{ "areacode":3, "number":1, "kind":"mobile" },
    {"areacode":4, "number":2, "kind":"mobile" }
]
WHERE id = 0
```

Output:

```
{
  "NumRowsUpdated" : 1
}
1 row returned
```

The result of the UPDATE statement can be verified as shown below.

```
select * from People
```

Output:

```
{
  "id" : 0,
  "info" : {
    "address" : {
      "city" : "Santa Cruz",
      "phones" : [[{
        "areacode" : 3,
        "kind" : "mobile",
        "number" : 1
      }, {
        "areacode" : 4,
        "kind" : "mobile",
        "number" : 2
      }], {
        "areacode" : 499,
        "kind" : "mobile",
        "number" : 33864368
      }, {
        "areacode" : 831,
        "kind" : "mobile",
        "number" : 5294368
      }, {
        "areacode" : 650,
        "kind" : "mobile",
        "number" : 3789021
      }, {
        "areacode" : 415,
        "kind" : "home",
        "number" : 6096010
      }, {
        "areacode" : 1,
        "kind" : "mobile",
        "number" : 1
      }, {
```



```

        "areacode" : 2,
        "kind" : "mobile",
        "number" : 2
    }],
    "state" : "CA"
},
"children" : {
    "Anna" : {
        "age" : 11,
        "friends" : ["Anna", "John", "Maria", "Ada", "Aris"],
        "school" : "school_1"
    },
    "Mary" : {
        "age" : 8,
        "friends" : ["Anna", "Mark", "Ada", "Aris"],
        "school" : "school_3"
    },
    "Ron" : {
        "age" : 3,
        "friends" : ["Julie", "Ada", "Aris"]
    }
},
"firstName" : "John",
"income" : 20000,
"lastName" : "Doe",
"profession" : "surfing instructor"
}
}

```

Example: Updating JSON collection tables

You can update data in the JSON collection tables using the UPDATE statement. The UPDATE operation works in the same way as fixed schema tables.

Consider a row from a JSON collection table created for a [shopping application](#).

```

{"contactPhone":"1617114988","address":{"Dropbox":"Presidency
College","city":"Kansas City","state":"Alabama","zip":95065},"cart":
[{"item":"A4 sheets","priceperunit":500,"quantity":2},{ "item":"Mobile
Holder","priceperunit":700,"quantity":1}], "email":"lorphil@usmail.com", "firstN
ame":"Lorenzo", "lastName":"Phil", "notify":"yes", "orders":
[{"EstDelivery":"2023-11-15", "item":"AG Novels
1", "orderID":"101200", "priceperunit":950, "status":"Preparing to dispatch"},
{"EstDelivery":"2023-11-01", "item":"Wall
paper", "orderID":"101200", "priceperunit":950, "status":"Transit"}]}

```

Example 7-40 Correct a few inadvertent errors in the shopper's data

Use [Update clauses](#) to correct a shopper's data as follows:

```

UPDATE storeAcct s
SET s.notify = "no",
REMOVE s.cart [$element.item = "A4 sheets"],
PUT s.address {"Block" : "C"},
SET s.orders[0].EstDelivery = "2023-11-17",

```

```
ADD s.cart 1 {"item":"A3 sheets", "priceperunit":600, "quantity":2}
WHERE s.contactPhone = "1617114988"
```

Explanation: In the above example, you update the shopper's record in the storeAcct table to correct a few inadvertent errors. This correction requires updates to various fields of the storeAcct table. The SET clause deactivates the notification setting in the shopper's data record. The REMOVE clause checks if any item field in the cart matches A4 sheets and deletes the corresponding element from the orders array. The PUT clause adds a new JSON field to indicate the landmark for delivery. The second SET clause accesses the deeply nested EstDelivery field and updates the estimated delivery date for the first item in the orders array. The ADD clause inserts a new element into the cart field to shortlist an additional item.

When you fetch the updated shopper's data, you get the following output:

```
{"contactPhone":"1617114988", "address":{"Block":"C", "Dropbox":"Presidency
College", "city":"Kansas City", "state":"Alabama", "zip":95065}, "cart":
[{"item":"Mobile Holder", "priceperunit":700, "quantity":1}, {"item":"A3
sheets", "priceperunit":600, "quantity":2}], "email":"lorphil@usmail.com", "firstN
ame":"Lorenzo", "lastName":"Phil", "notify":"no", "orders":
[{"EstDelivery":"2023-11-17", "item":"AG Novels
1", "priceperunit":950, "status":"Preparing to dispatch"},
{"EstDelivery":"2023-11-01", "item":"Wall
paper", "priceperunit":950, "status":"Transit"}]}
```

Example: Updating TTL

This example demonstrates an update of the expiration time of a row. For creating the table and inserting data, see [Example: Updating Rows](#).

Example 7-41 Update the TTL value for a table

Let's assume that the People table was created with a TTL value of 10 hours and a row with id 5 was inserted at time 2017-06-01T10:05:30.0. No explicit TTL was given at insertion time, so the expiration time computed at that time is 2017-06-01T21:00:00.0. Finally, let's assume that the following update statement is executed at time 2017-06-01T12:35:30.0 (2.5 hours after insertion)

```
UPDATE People $p
SET TTL remaining_hours($p) + 3 hours
WHERE id = 5
```

The above statement extends the life of a row by 3 hours. Specifically, the remaining_hours function returns the number of full hours remaining until the expiration time of the row. See [remaining_hours function](#) function. In this example, this number is 8. So, the new TTL value is 8+3 = 11, and the expiration time of the row will be set to 2017-06-02:T08:00:00.0.

Notice the use of the '\$' character in naming the table alias for People. This is required so that the table alias acts as a row variable (a variable ranging over the rows of the table) and as a result it can be passed as the argument to the remaining_hours function (if the \$ were not used, then calling remaining_hours(p) would return an error, because p is interpreted as a reference to a top-level table column with name "p").

Example: Updating IDENTITY defined as GENERATED ALWAYS

Create a table with primary key as an IDENTITY column defined as GENERATED ALWAYS AS IDENTITY and insert data to the table.

```
CREATE TABLE Test_sqlUpdateAlways (  
    idValue INTEGER GENERATED ALWAYS AS IDENTITY,  
    name STRING,  
    PRIMARY KEY(idValue))
```

```
INSERT INTO Test_sqlUpdateAlways VALUES (DEFAULT, 'joe')
```

```
INSERT INTO Test_sqlUpdateAlways VALUES (DEFAULT, 'jasmine')
```

The Test-sqlUpdateAlways table will have the following rows:

```
1, 'joe'  
2, 'jasmine'
```

Example 7-42 Error scenario - Update IDENTITY column defined as GENERATED ALWAYS

idValue

```
UPDATE Test_sqlUpdateAlways SET idValue = 10 WHERE name=joe
```

The UPDATE statement above will raise an exception saying that a user cannot set a value for an IDENTITY column that is defined as GENERATED ALWAYS. An IDENTITY column that is defined as GENERATED ALWAYS cannot be updated. Only the IDENTITY column that is defined as GENERATED BY DEFAULT can be updated.

To resolve this exception and be able to update the IDENTITY column value, you must alter the IDENTITY column and change the property of the IDENTITY column to GENERATED BY DEFAULT. But there can be implications to the existing data. For more information on how to alter an IDENTITY column see, [Altering an IDENTITY Column](#).

Example: Updating IDENTITY defined as GENERATED BY DEFAULT

Create a table with primary key as an IDENTITY column defined as GENERATED BY DEFAULT and insert data to the table.

```
CREATE TABLE Test_sqlUpdateByDefault (  
    idValue INTEGER GENERATED BY DEFAULT AS IDENTITY,  
    acctNum LONG,
```

```
name STRING,  
primary key(acctNum))
```

```
INSERT INTO Test_sqlUpdateByDefault VALUES (DEFAULT, 123456, 'joe')
```

```
INSERT INTO Test_sqlUpdateByDefault VALUES (400, 23456, 'sam')
```

```
INSERT INTO Test_sqlUpdateByDefault VALUES (500, 34567, 'carl')
```

Table Test-sqlUpdateByDefault will have the following rows:

```
1, 123456, 'joe'  
400, 23456, 'jasmine'  
500, 34567, 'carl'
```

Example 7-43 Update IDENTITY column defined as GENERATED BY DEFAULT

```
UPDATE Test_sqlUpdateByDefault  
SET idValue = 100  
WHERE acctNum = 123456
```

The UPDATE statement above replaces the row (1, 123456, 'joe') with (100, 123456, 'joe') in the database.

JSON Collection Table Example

This section describes a sample JSON collection table created for a shopping application.

Table DDL:

```
CREATE TABLE IF NOT EXISTS storeAcct (contactPhone STRING, PRIMARY  
KEY(SHARD(contactPhone)))  
AS JSON COLLECTION
```

This table is a collection of documents with the shopper's contactPhone as the primary key. The rows represent individual shopper's records. The individual rows need not include the same fields in the document. The shopper's preferences such as name, address, email, notify, and so forth are stored as top-level fields in the document. The documents can include any number of JSON fields such as wishlist, cart, and orders that contain shopping-related information.

The JSON array wishlist contains the items wishlisted by the shoppers. Each element of this array includes nested JSON fields such as the item and priceperunit to store the product name and price details of the wishlisted item.

The JSON array cart contains the products that the shopper intends to purchase. Each element of this array includes nested JSON fields such as item, quantity, and priceperunit to store the product name, number of units, and price of each unit.

The JSON array orders contains the products that the shopper has purchased. Each element of this array includes nested JSON fields such as the orderID,

item, priceperunit, EstDelivery, and status to store the order number, product name, price of each unit, estimated date of delivery for the product, and status of the order.

The following code inserts data into the shopping application table. You can use this data to follow along with the examples explained in the topics.

```
insert into storeAcct(contactPhone, firstName, lastName, address, cart)
values("1817113382", "Adam", "Smith", {"street" : "Tex Ave", "number" : 401,
"city" : "Houston", "state" : "TX", "zip" : 95085}, [{"item" : "handbag",
"quantity" : 1, "priceperunit" : 350}, {"item" : "Lego", "quantity" : 1,
"priceperunit" : 5500}]) RETURNING *
```

```
insert into storeAcct(contactPhone, firstName, lastName, gender, address,
notify, cart, wishlist) values("1917113999", "Sharon", "Willard", "F",
{"street" : "Maine", "number" : 501, "city" : "San Jose", "state" : "San
Francisco", "zip" : 95095}, "yes", [{"item" : "wallet", "quantity" : 2,
"priceperunit" : 950}, {"item" : "wall art", "quantity" : 1, "priceperunit" :
9500}], [{"item" : "Tshirt", "priceperunit" : 500}, {"item" : "Jenga",
"priceperunit" : 850}]) RETURNING *
```

```
insert into storeAcct(contactPhone, firstName, lastName, address, notify,
cart, orders) values("1617114988", "Lorenzo", "Phil", {"Dropbox" :
"Presidency College", "city" : "Kansas City", "state" : "Alabama", "zip" :
95065}, "yes", [{"item" : "A4 sheets", "quantity" : 2, "priceperunit" : 500},
{"item" : "Mobile Holder", "quantity" : 1, "priceperunit" : 700}], [{"orderId" : "101200", "item" : "AG Novels 1", "EstDelivery" : "2023-11-15",
"priceperunit" : 950, "status" : "Preparing to dispatch"}, {"orderId" :
"101200", "item" : "Wallpaper", "EstDelivery" : "2023-11-01",
"priceperunit" : 950, "status" : "Transit"}]) RETURNING *
```

```
insert into storeAcct(contactPhone, firstName, lastName, address, cart,
orders) values("1517113582", "Dierdre", "Amador", {"street" : "Tex Ave",
"number" : 651, "city" : "Houston", "state" : "TX", "zip" : 95085}, NULL,
[{"orderId" : "201200", "item" : "handbag", "EstDelivery" : "2023-11-01",
"priceperunit" : 350}, {"orderId" : "201201", "item" : "Lego", "EstDelivery" :
"2023-11-01", "priceperunit" : 5500}]) RETURNING *
```

```
insert into storeAcct(contactPhone, firstName, lastName, address, notify,
cart, orders) values("1417114488", "Doris", "Martin", {"Dropbox" :
"Presidency College", "city" : "Kansas City", "state" : "Alabama", "zip" :
95065}, "yes", [{"item" : "Notebooks", "quantity" : 2, "priceperunit" : 50},
{"item" : "Pens", "quantity" : 2, "priceperunit" : 50}], [{"orderId" :
"301200", "item" : "Laptop Bag", "EstDelivery" : "2023-11-15",
"priceperunit" : 1950, "status" : "Preparing to dispatch"}, {"orderId" :
"301200", "item" : "Mouse", "EstDelivery" : "2023-11-02", "priceperunit" :
950, "status" : "Transit"}]) RETURNING *
```

8

Indexes

This chapter describes indexes and how to create and manage indexes in Oracle NoSQL Database.

This chapter contains the following topics:

- [About Indexes](#)
- [Classification of Indexes](#)
- [CREATE INDEX Statement](#)
- [SHOW INDEXES Statement](#)
- [DESCRIBE INDEX Statement](#)
- [DROP INDEX Statement](#)
- [Appendix](#)

To follow along with the examples in the topics, create the UserInfo table as described in the [Appendix](#) section.

About Indexes

An index is a database structure that enables you to retrieve data from database tables efficiently.

Indexes provide fast access to the rows of a table when the key(s) you are searching on is contained in the index.

An index is defined by its name, the name of the table that it indexes, and a list of one or more index paths that specify which table columns or nested fields are indexed.

An index is an ordered map in which each row of the data is called an entry.

An index can be created on atomic data types, arrays, maps, JSON, and GeoJSON data.

An index can store the following special values:

- NULL
- EMPTY
- json null (It is applicable only for JSON indexes)

Example 8-1 Indexes Example

The following is an example of creating an index. The index is created on the age field present in the info JSON field in the UserInfo table.

```
CREATE INDEX indexdemo1 ON UserInfo(info.age AS ANYATOMIC)
```

CREATE INDEX Statement

The create index statement generates a new index on the specified columns in the given table.

Syntax

```
create_index_statement ::=
    CREATE INDEX [IF NOT EXISTS] index_name
    ON table_name "(" path_list ")" [WITH NO NULLS][WITH UNIQUE KEYS PER ROW]
    [comment]

index_name ::= id
path_list ::= index_path ("," index_path)*
index_path ::=
    name_path [path_type] |
    multikey_path_prefix [.name_path] [path_type]
name_path ::= field_name ("." field_name)*
field_name ::= id | DSTRING
multikey_path_prefix ::=
    field_name ( ( "." field_name ) | ( "[" "]" ) | ( "." VALUES "(" )" ) ) *
    ( ( "[" "]" ) | ( "." VALUES "(" )" ) | ( "." KEYS "(" )" ) )
path_type ::= AS
    (INTEGER | LONG | DOUBLE | STRING |
    BOOLEAN | NUMBER | ANYATOMIC | POINT | GEOMETRY)
```

Semantics

The index name is unique to a table. If an index with the same name already exists in a table, then the statement will fail and report an error. For example, you can have only one index named `idx_income1` in the `UserInfo` table.

The index name is specific to a table. You can use the same index name in multiple tables. For example, you can use the same index name `idx_income1` in the `UserInfo` and `Users3` tables.

The index specification is unique to a table. If an index with the same specification already exists in a table, then the statement will fail and report an error.

For example, if you have the following `idx_income1` index on `UserInfo` table,

```
CREATE INDEX idx_income1 ON UserInfo (info.income AS ANYATOMIC)
```

then the following statement will throw an error that the `idx_income2` index is a duplicate of an existing index with another name. In this case, even though the index names are different, the index specifications are the same.

```
CREATE INDEX idx_income2 ON UserInfo(info.income AS ANYATOMIC)
```

If the optional `IF NOT EXISTS` clause is specified in the `CREATE INDEX` statement, and if an index with the same name exists, then the statement will not execute and will not report an error.

If the optional `WITH NO NULLS` clause is specified in the `CREATE INDEX` statement, then the rows with `NULL` and/or `EMPTY` values on the indexed fields will not be indexed.

The indexes that are created with the WITH NO NULLS clause may be useful when the data contain a lot of NULL and/or EMPTY values on the indexed fields. It will reduce the time and space overhead during indexing. However, the use of such indexes by queries is restricted. For more information, see [Using Indexes for Query Optimization](#) section.

If the optional WITH UNIQUE KEYS PER ROW clause is used, then there will not be any duplicates among the index keys generated from a row. This property applies to multikey indexes only and is used in optimizing queries that perform unnesting. You could write an efficient query to use this index. The use of such an index by any query would yield fewer results from the FROM clause than if the index was not used.

```
CREATE INDEX idx_phones ON UserInfo(info.phones[].number AS INTEGER)
WITH UNIQUE KEYS PER ROW
```

If the optional COMMENT is specified, then this becomes part of the index metadata and is not interpreted. The "comment" will be displayed in the output of the DESCRIBE statement.

If JSON data is indexed, you must specify a data type using the AS keyword next to every index path into the JSON data. For all other typed data, you should not specify the data type, as the data type will be inferred from the table schema.

The index entries are automatically updated when rows are inserted, deleted, or updated in the specified table.

Note: The maximum number of index keys generated per row is 10000. An `IllegalArgumentException` will be raised during indexing if the number of index keys generated per row exceeds 10000.

Classification of Indexes

Indexes can be classified based on fields, schema, entries, or a combination of them. Each one of these is described below.

- Fields
 - Single Field Index
 - Composite Index
- Schema
 - Fixed Schema Index
 - Schema-less Index (JSON Index)
- Entries
 - Simple Index
 - Multikey Index

The following classification is made based on the number of fields that are provided while creating an index.

Single Field Index

An index is called a single field index if it is created on only one field of a table.

Composite Index

An index is called a composite index if it is created on more than one field of a table.

The following classification is made based on the schema type of the fields that are indexed.

Fixed Schema Index

An index is called a fixed schema index if all the fields that are indexed are strongly typed data. For more information on strongly typed data, see [Wildcard Data Types](#) section.

Schema-less Index (JSON Index)

An index is called a JSON index if at least one of the fields is JSON data or fields inside JSON data.

The following classification is made based on the number of index entries created for each row of data in the table when evaluating an index.


Simple Index

An index is called a simple index if for each row of data in the table, there is one entry created in the index.

Multikey Index

An index is called a multikey index if for each row of data in the table, there are multiple entries created in the index.

Figure 8-1 Index Classification



Single Field Index

An index is called a single field index if it is created on only one field of a table.

Example 8-2 Single Field Index

The following is an example of a single field index. The index is created on the city field present in the address record in the info JSON field in the UserInfo table.

```
CREATE INDEX singlefieldindex1 ON UserInfo (  
    info.address.city AS ANYATOMIC  
)
```

Composite Index

An index is called a composite index if it is created on more than one field of a table.

Example 8-3 Composite Index

The following is an example of a composite index. The index is created on the state and city fields present in the address record in the info JSON field in the UserInfo table.

```
CREATE INDEX compositeindex1 ON UserInfo (  
    info.address.state AS ANYATOMIC,  
    info.address.city AS ANYATOMIC  
)
```

Fixed Schema Index

An index can be created on a field with fixed schema data.

Example 8-4 Fixed Schema Index

The following is an example of a fixed schema index. The index is created on the uname field having integer data type in the UsersInfo table.

```
CREATE INDEX fixedschemaindex1 ON UserInfo (uname)
```

JSON Index

An index is called a JSON index if at least one of the fields is inside JSON data.

As JSON is schema-less, the data type of an indexed JSON field may be different across rows. When creating an index on JSON fields, if you are unsure what data type to expect for the JSON field, you may use the anyAtomic data type. Alternatively, you can specify one of the Oracle NoSQL Database atomic data types. You do that by declaring a data type using the AS keyword next to every index path into the JSON field.

Example 8-5 JSON Index

The following is an example of a JSON index. The index is created on the income field present in the info JSON field in the UserInfo table. Notice that you provide a data type for the income field while creating the index.

```
CREATE INDEX jsonindex1 ON UserInfo (  
    info.income AS INTEGER  
)
```

The creation of a JSON index will fail if the associated table contains any rows with data that violate the declared data type. Similarly, after creating a JSON index, an insert/update operation will fail if the new row does not conform to the declared data type in the JSON index.

For example, the jsonindex1 index will be created only if the income field in all the rows of the UserInfo table is of integer data type, if not the jsonindex1 index creation will fail. Similarly, after creating the jsonindex1 index on the UserInfo table, you can insert only rows in which the income field is of integer data type. For example, If you try inserting a row in which the income field is of string data type, the insert statement will fail.

Declaring a JSON index path as anyAtomic has the advantage of allowing the indexed JSON field to have values of various data types. The index entries are sorted in ascending order. When these values are stored in the index, they are sorted as follows:

- Numbers

- String
- boolean

However, this advantage is offset by space and CPU costs. It is because numeric values of any kind in the indexed field will be cast to Number before being stored in the index. This cast takes CPU time, and the resulting storage for the number will be larger than the original storage for the number.

Example 8-6 JSON Index Example

The following is an example of a JSON index. The index is created on the `street` field present in the `address` field in the `info` JSON document in the `UserInfo` table. Notice that you provide `anyAtomic` data type for the `street` field while creating the index.

```
CREATE INDEX jsonindex2 ON UserInfo (  
    info.address.street AS ANYATOMIC  
)
```

Simple Index

An index is called a simple index if, for each row of data in the table, there is one entry created in the index. The index will return a single value that is of atomic data type or any special value (SQL NULL, JSON NULL, EMPTY). Essentially, the index paths of a simple index must not have `.keys()`, or `.values()`, or `[]` steps.

Example 8-7 Simple Index

The following is an example of a simple index. The index is created on the `income` and `age` fields present in the `info` JSON field in the `UserInfo` table.

```
CREATE INDEX simpleindex1 ON UserInfo (  
    info.income AS ANYATOMIC,  
    info.age AS ANYATOMIC  
)
```

A simple index path must not include an array since arrays will render multiple index values. For example, `info.connections[]` returns the complete set of elements in the `info.connections` array. It is not a simple index.

Similarly, a simple index path must not have a map with `.keys()` or `.values()`. For example, `info.expenses.keys()` returns the complete set of keys in the `expenses` map. It is not a simple index path. However, `info.expenses.books` is a simple index path. Because even though `info.expenses` is a map, `info.expenses.books` return an atomic value.

Example 8-8 Simple Index Example

The following is an example of a simple index created on a JSON document in a JSON field. The index is created on the `books` item of the `expenses` JSON document in the `info` JSON field in the `UserInfo` table.

```
CREATE INDEX simpleindex2 ON UserInfo (  
    info.expenses.books AS ANYATOMIC  
)
```

If the evaluation of a simple index path returns an empty result, the special value `EMPTY` is used as an index entry. In the above example, If there is no books entry in the expenses JSON document, or if there is no expenses JSON document, then the special value `EMPTY` is indexed.

Multikey Index

An index is called a multikey index if for each row of data in the table, there are multiple entries created in the index. In a multikey index there is at least one index path that uses `.keys()`, `.values()`, or `[]` steps. Any such index path will be called a multikey index path.

In a multikey index, for each table row, index entries are created on all the elements in arrays or entries in maps that are being indexed. If the evaluation returns an empty result, the special value `EMPTY` is used as the index entry. Any duplicate index entries are then eliminated.

Example 8-9 Multikey Index

The following is an example of a multikey index. The index is created on the `connections[]` array in the `UserInfo` table. Here, all the elements in the `connections[]` array in each row of the `UserInfo` table will be indexed.

```
CREATE INDEX multikeyindex1 ON UserInfo (  
    info.connections[] AS ANYATOMIC  
)
```

Nested Multikey Index

An index is a nested multikey index if it is created on a field that is present inside an array which in turn is present inside another array.

Example 8-10 Nested Multikey Index

The following is an example of a nested multikey index where the field is present in an array that is present inside another array. The index is created on the `issuedby` field in the `vpass` array in the `vehicles` array in the `info` JSON of the `UserInfo` table.

```
CREATE INDEX multikeyindex2 ON UserInfo (  
    info.vehicles[].vpass[].issuedby AS ANYATOMIC  
)
```

Composite Multikey Index

An index is called a composite multikey index if it is created on more than one field, and at least one of those fields is multikey.

A composite multikey index may have a combination of multikey index paths and simple index paths.

Example 8-11 Composite Multikey Index

The following is an example of a composite multikey index having one multikey index path and one simple index path. The index is created on the `income` field and `area` field in the `info` JSON column of the `UserInfo` table.

```
CREATE INDEX multikeyindex3 ON UserInfo (
    info.income AS ANYATOMIC,
    info.address.phones[].area AS ANYATOMIC
)
```

A composite multikey index may have more than one multikey index path. Some of the possibilities of composite multikey indexes are given below.

1. You can use both the `.keys()` and `.values()` steps of a map or JSON together.

The following is an example of a composite multikey index in which both the `.keys()` and `.values()` steps of a JSON document are used together. The index is created on the keys and values of the `expenses` JSON document in the `UserInfo` table.

```
CREATE INDEX multikeyindex4 ON UserInfo (
    info.expenses.keys(),
    info.expenses.values() AS ANYATOMIC
)
```

2. You can use multiple fields of an array/record/maps-viewed-as-records together. However, the restriction here is that a field cannot be treated as both array and map within a single `CREATE INDEX` statement.

The following is an example of a composite multikey index in which multiple fields of an array are used together. The index is created on the `passid` and `issuedby` fields in the `vpass` array in the `UserInfo` table.

```
CREATE INDEX nestedindex1 ON UserInfo (
    info.vehicles[].vpass[].passid AS ANYATOMIC,
    info.vehicles[].vpass[].issuedby AS ANYATOMIC
)
```

Figure 8-2 Composite Multikey Index



3. You can use fields of an array/record/maps-viewed-as-records, as well as the fields present inside the inner arrays of those array/record/maps-viewed-as-records together. However, the restriction here is that the immediate parent array of all such fields should be the same.

The following is an example of a composite multikey index in which fields of an array and fields of an inner array are used together. The index is created on the vid field in the vehicles array and the passid field in the vpass array which is in the vehicles array in the UserInfo table.

```
CREATE INDEX nestedindex2 ON UserInfo (
    info.vehicles[].vid AS ANYATOMIC,
    info.vehicles[].vpass[].passid AS ANYATOMIC
)
```

Figure 8-3 Composite Multikey Index



The following is an example of a composite multikey index in which fields of an array and fields of an inner array are used together, however, notice that the immediate parent array of the two fields are not the same. The following is an invalid index creation statement. The index is being created on the passid and serviceid fields in the vehicles array in the UserInfo table. Note that within the vehicles array, passid is inside the vpass array, and serviceid is inside the vservice array.

```
CREATE INDEX invalidindex1 ON UserInfo (
    info.vehicles[].vpass[].passid AS ANYATOMIC,
    info.vehicles[].vservice[].serviceid AS ANYATOMIC
)
```

Figure 8-4 Composite Multikey Index (invalid)



Specifications & Restrictions on Multikey Indexes

The following specifications & restrictions apply to multikey index paths:

- You cannot provide predicate or boundary expressions for `.keys()`, `.values()`, and `[]` steps. For more information on predicate and boundary expressions, see the [Map-Filter Step Expressions](#) and [Array-Slice Step Expressions](#) sections respectively.
- When a multikey index path is evaluated on a table row, it must return zero or more atomic values. If no value is returned, then the special value `EMPTY` is used.
- No data type declaration is allowed for `.keys()` step, as the `keys()` are by default string data type.
- You can provide `.keys()` steps only as the last step in an index path.
- For Example, `info.expenses.keys()` is a valid index path whereas `info.expenses.keys().books` is an invalid index path.
- You can provide `.keys()` and `.values()` only after the fields that are maps in strongly typed data.
- You can provide `.keys()` and `.values()` steps after a field that you expect to be atomic values or JSON documents inside JSON data.
 - If the fields are atomic values, then the special value `EMPTY` will be returned.
 - If the fields are json documents, then the corresponding keys or values will be returned.
- You should provide `[]` steps after those fields that are arrays in strongly typed data.
- You should provide `[]` steps after those fields that you expect to be arrays inside JSON data.
- If you do not provide `[]` steps after those fields that you expect to be arrays inside JSON data, then those fields cannot be an array.
 - If the rows in the existing table contain any array inside the JSON data, then the index creation will fail.
 - However, if there are no rows in the existing table that contain an array inside the JSON data, then the index creation will be successful. However, you will not be able to insert a new row with an array in that JSON data in the existing table.
 - For example, in the following `badindex1`, the `[]` step is not used with the `vehicles` array. There can be two scenarios here:

```
CREATE INDEX badindex1 ON UserInfo(info.vehicles.vid AS ANYATOMIC)
```

- * If there are existing rows of data in the `UserInfo` table where the `vehicles` field happens to be an array, the index creation will fail.
- * If there are no existing rows of data with the `vehicles` array in the `UserInfo` table, the index will be created successfully. However, if you attempt to insert a new row of data where the `vehicles` attribute is an array, the row insertion will fail.
- You cannot provide `[]` steps after the non-array fields in strongly typed data.
- You can provide `[]` steps after the atomic values and JSON documents in JSON data. If you do so, then that atomic value or JSON document will be used.
- If an array contains another array directly, then the corresponding number of `[]` steps are required to index the elements of the inner arrays. For example, if there are two inner arrays inside `vehicles`, then `info.vehicles[][].vid` should be used.

Index on JSON Collection Tables

You can index the fields in a JSON collection table. You must specify the name of the indexed element and ANYATOMIC for the type definition, or, for strongly typed indexes, you can specify the JSON type of the fields being indexed.

Indexing the fields in the JSON collection table is similar to creating [JSON indexes](#). You must specify the path expression to the field. If you are indexing a top-level JSON field in the document, the field name is its path expression. If the element is deeply nested in a JSON object, you specify the complete path name. In either case, the data type for every index must be specified and it is recommended to use ANYATOMIC in the JSON collection tables.

Consider the JSON collection table created for a [shopping application](#).

Example 8-12 Create an index on the JSON collection table

```
create index myindex on usersJSON(notify as ANYATOMIC)
```

In the statement above, you create an untyped index on the notify field of the shopper's data.

Example 8-13 Create a composite index on the JSON collection table

```
create index idx_ntfy_cty on storeAcct (address.city as ANYATOMIC, notify as ANYATOMIC)
```

The notify field is a top-level field and the city field is nested in the address field of the storeAcct table. In this statement, you create a composite index using these two fields.

Note

If you are creating an index on a nested JSON field, the field must be present in all the rows of the table. Otherwise, an error is displayed.

Example 8-14 Create a strongly typed index on the JSON collection table

```
create index myindex2 on usersJSON(notify as string)
```

In the statement above, you create a string index on the top-level notify field of the shopper's data.

Note that the creation of a strongly typed index will fail if the table includes any rows with data that violate the declared data type. Also, after successfully creating this index, you can only insert string data into the notify field. You can use a strongly typed index to act as both an index and a type constraint on a JSON field.

You can drop the indexes on the JSON collection table using the DROP INDEX Statement. For details, see [DROP INDEX Statement](#).

Indexes on Functions

You can create indexes on the values of one or more SQL built-in functions.

List of functions that can be indexed:

The following subset of the [Built-in Functions](#) can be indexed.

Functions on Timestamps:

- year
- month
- day
- hour
- minute
- second
- millisecond
- microsecond
- nanosecond
- week

Functions on Strings:

- length
- replace
- reverse
- substring
- trim
- ltrim
- rtrim
- lower
- upper

Functions on Rows:

- modification_time
- expiration_time
- expiration_time_millis
- row_storage_size

See [Built-in Functions](#) for more details on what a built-in function is and how to use these functions.

Examples of creating indexes on functions

You can create indexes on the values of one or more SQL built-in functions.

Create Index on row-property functions

A row-property function expects a row variable as its sole argument in a DML statement. You can create an index on these row property functions.

- modification_time
- expiration_time
- expiration_time_millis

- `row_storage_size`

In a `CREATE INDEX` statement, you must provide these functions without any argument. The row you are indexing is implicitly considered as the input to the function.

Example:

Create an index which indexes the rows of the `Users` table by its latest modification time:

```
CREATE INDEX idx_modtime ON users(modification_time())
```

This index will be used in a query, which has `modification_time` as the filter condition.

```
SELECT * FROM Users $u WHERE  
modification_time($u) > "2022-08-01T10:45:00"
```

This query returns all the rows whose most recent modification time is after 2022-08-01T10:45:00. It uses the `idx_modtime` index defined above. You can verify this by viewing the query plan using the `show query` command.

More examples of creating indexes on functions

If you are creating an index on a built-in function that is not a row property function, then the first argument of the function must be an index path. This is followed by a path type if the path points to JSON data as shown in the example below. Some string functions need more than one argument and in that case, all the arguments other than index path and path type should be constant literal only.

Note

The path type is not needed if it is a non-JSON as the datatype is the same as the definition in the table. For JSON data, the type `ANYATOMIC` can be used to cover all valid types in JSON.

Characteristics of the built-in functions that can be indexed:

- All of the built-in functions that can be indexed expect at most one item as their first argument.
- The time-related functions such as `YEAR`, `MONTH` etc will throw an exception if their input is a sequence with more than 1 item, and the string related functions such as `length`, `reverse` and so on will return `NULL` in this case.
- If the input index path has a multi-key, then the function will be evaluated separately for each value returned by the multi-key path. The resulting sequence of items will be combined with the values of the other index paths (if any) in the index definition and index entries will be created.

Download the script `acctstream_loaddata.sql` and run it as shown below. This script creates the `stream_acct` table used in the example below and loads data into the table.

Start your `KVSTORE` or `KVlite` and open the `SQL.shell`.

```
java -jar lib/kvstore.jar kvlite -secure-config disable  
java -jar lib/sql.jar -helper-hosts localhost:5000 -store kvstore
```

Using the load command, run the script.

```
load -file acctstream_loaddata.sql
```

Example 1: Create an index in the stream_acct table to index users by the first three letters of their country of residence.

```
CREATE INDEX idx_country ON stream_acct(substring(acct_data.country as
string, 0,3))
```

- In the example above, you declare the **country** field of the **acct_data** JSON as a string.
- The substring function is used to extract the first three letters of the country and use it as the index key.

Example 2: Using a multi-key index path

In the following example, you index the users in the stream_acct table by the id of the shows they watch and the year and month of the dates when the show was watched.

```
CREATE INDEX idx_showid_year_month ON
stream_acct(acct_data.contentStreamed[].showId AS INTEGER,
substring(acct_data.contentStreamed[].seriesInfo[].episodes[].date AS
STRING,0, 4),
substring(acct_data.contentStreamed[].seriesInfo[].episodes[].date AS
STRING,5, 2))
```

An example of a query using this index is shown below. The query counts the number of users who watched any episode of show **16** in the year **2022**.

```
SELECT count(*) FROM stream_acct s1 WHERE EXISTS
s1.acct_data.contentStreamed[$element.showId = 16].seriesInfo.
episodes[substring($element.date, 0, 4) = "2022"]
```

This query will use the index idx_showid_year_month. You can verify this by viewing the query plan using the show query command.

```
show query SELECT count(*) FROM stream_acct s1 WHERE EXISTS
> s1.acct_data.contentStreamed[$element.showId =
16].seriesInfo.episodes[substring($element.date, 0, 4) = "2022"]
```

```
{
  "iterator kind" : "GROUP",
  "input variable" : "$gb-1",
  "input iterator" :
  {
    "iterator kind" : "RECEIVE",
    "distribution kind" : "ALL_SHARDS",
    "distinct by fields at positions" : [ 1 ],
    "input iterator" :
    {
      "iterator kind" : "SELECT",
      "FROM" :
      {
```

```

        "iterator kind" : "TABLE",
        "target table" : "stream_acct",
        "row variable" : "$s1",
        "index used" : "idx_showid_year_month",
        "covering index" : true,
        "index row variable" : "$s1_idx",
        "index scans" : [
            {
                "equality conditions" :
{"acct_data.contentStreamed[].showId":16,"substring#acct_data.contentStreamed[
].seriesInfo[].episodes[].date@,0,4":"2022"},
                "range conditions" : {}
            }
        ]
    },
    "FROM variable" : "$s1_idx",
    "SELECT expressions" : [
        {
            "field name" : "Column_1",
            "field expression" :
            {
                "iterator kind" : "CONST",
                "value" : 1
            }
        },
        {
            "field name" : "acct_id_gen",
            "field expression" :
            {
                "iterator kind" : "FIELD_STEP",
                "field name" : "#acct_id",
                "input iterator" :
                {
                    "iterator kind" : "VAR_REF",
                    "variable" : "$s1_idx"
                }
            }
        }
    ]
}
},
"grouping expressions" : [

],
"aggregate functions" : [
    {
        "iterator kind" : "FUNC_COUNT_STAR"
    }
]
}

```

SHOW INDEXES Statement

The SHOW INDEXES statement provides the list of indexes present on the specified table.

Syntax

```
show_indexes_statement ::=  
    SHOW [AS JSON] INDEXES ON table_name
```

Semantics

If you want the output to be in JSON format, you can specify the optional AS JSON.

Example 8-15 Show Indexes

The following statement lists the indexes present on the UserInfo table.

```
SHOW INDEXES ON UserInfo
```

Output

```
indexes  
  idx_phones  
  idx_income1
```

Example 8-16 Show Indexes

The following statement lists the indexes present on the UserInfo table in JSON format.

```
SHOW AS JSON INDEXES ON UserInfo
```

Output

```
{"indexes"  
  : ["idx_phones","idx_income1"]}
```

DESCRIBE INDEX Statement

The DESCRIBE INDEX statement defines the specified index on a table.

Syntax

```
describe_index_statement ::=  
    (DESCRIBE | DESC) [AS JSON] INDEX index_name ON table\_name
```

Semantics

If you want the output to be in JSON format, you can specify the optional AS JSON.

The description for the index contains the following information:

- Name of the table on which the index is defined.
- Name of the index.
- Type of index. Whether the index is primary index or secondary index.

- Whether the index is multikey? If the index is multikey then 'Y' is displayed. Otherwise, 'N' is displayed.
- List of fields on which the index is defined.
- The declared type of the index.
- Description of the index.

Example 8-17

The following statement provides information about the index `idx_income1` on the `UserInfo` table.

```
DESCRIBE AS JSON INDEX idx_income1 ON UserInfo
```

Output

```
{
  "name" : "idx_income1",
  "type" : "secondary",
  "fields" : [{
    "path" : "info.income",
    "type" : "AnyAtomic"
  }],
  "withNoNulls" : false,
  "withUniqueKeysPerRow" : false
}
```

Example 8-18

The following statement provides information about the index `idx_phones` on the `UserInfo` table.

```
DESCRIBE AS JSON INDEX idx_phones ON UserInfo
```

Output

```
{
  "name" : "idx_phones",
  "type" : "secondary",
  "fields" : [{
    "path" : "info.phones[].number",
    "type" : "Integer"
  }],
  "withNoNulls" : false,
  "withUniqueKeysPerRow" : true
}
```

Example 8-19

The following statement provides information about the index `idx_modtime` on the `users` table.

```
DESCRIBE AS JSON INDEX idx_modtime ON users
```

Output

```
{
  "name" : "idx_modtime",
  "type" : "secondary",
  "fields" : [{
    "function" : "modification_time",
    "type" : "Timestamp(3)"
  }],
  "withNoNulls" : false,
  "withUniqueKeysPerRow" : false
}
```

Example 8-20

The following statement provides information about the index `idx_showid_year_month` on the `stream_acct` table.

```
DESCRIBE AS JSON INDEX idx_showid_year_month ON stream_acct
```

Output:

```
{
  "name" : "idx_showid_year_month",
  "type" : "secondary",
  "fields" : [{
    "path" : "acct_data.contentStreamed[].showId",
    "type" : "Integer"
  },
  {
    "function" : "substring",
    "arguments" : "acct_data.contentStreamed[].seriesInfo[].episodes[].date
AS String,0,4",
    "type" : "String"
  },
  {
    "function" : "substring",
    "arguments" : "acct_data.contentStreamed[].seriesInfo[].episodes[].date
AS String,5,2",
    "type" : "String"
  }
  ],
  "withNoNulls" : false,
  "withUniqueKeysPerRow" : false
}
```

DROP INDEX Statement

The DROP INDEX removes the specified index from the database.

Syntax

```
drop_index_statement ::=  
    DROP INDEX [IF EXISTS] index_name ON table\_name
```

Semantics

If an index with the given name does not exist, then the statement fails, and an error is reported.

If the optional IF EXISTS clause is used in the DROP INDEX statement, and if an index with the same name does not exist, then the statement will not execute, and no error is reported.

Example 8-21 Drop index

```
DROP INDEX idx_income1 ON UserInfo
```

Output:

Statement completed successfully

Appendix

The following code creates the UserInfo table.

```
CREATE TABLE UserInfo (id INTEGER, uname STRING, info JSON, PRIMARY KEY(id))
```

The following code populates the UsersInfo table with sample rows.

```
INSERT INTO UserInfo VALUES (  
    1001,  
    "Peter",  
    {  
        "age":42,  
        "income":65000,  
        "address": {  
            "street":"Lane-8",  
            "city":"Boston",  
            "state":"MA",  
            "phones":[  
                {"area":415,"number":91237468,"kind":"work"},  
                {"area":null,"number":95213607,"kind":"home"}  
            ]  
        },  
        "vehicles" : [  
            {  
                "vid":72132,  
                "vtype":"car",  
                "vpass":[  
                    {"passid":396457,"issuedby":"BPD"},  
                    {"passid":312358,"issuedby":"NYPD"}  
                ],  
            },  
        ],  
    )
```



```

        "vservice":[
            {"serviceid":20001,"servicedate":null}
        ],
        {
            "vid":78344,
            "vtype":"bike",
            "vpass":[
                {"passid":396241,"issuedby":"BPD"}
            ]
        }
    ],
    "expenses":{"housing":1000,"clothes":230,"books":20},
    "connections":[100,20,20,10,20]
}
)

INSERT INTO UserInfo VALUES (
    1002,
    "Ram",
    {
        "age":35,
        "income":null,
        "address":{"
            "street":"Hosur Road",
            "city":"Bengaluru",
            "state":"KA",
            "phones":[
                {"area":080,"number":2653457,"kind":"work"},
                {"area":080,"number":2659753,"kind":"home"}
            ]
        },
        "vehicles":null,
        "expenses":{"housing":1000,"travel":300},
        "connections":[ ]
    }
)

INSERT INTO UserInfo VALUES (
    1003,
    "Alice",
    {
        "income":20000,
        "address":{"
            "street":"Fremont Rd",
            "city":"San Jose",
            "state":"CA",
            "phones":[ ]
        },
        "expenses":null,
        "connections":null
    }
)

INSERT INTO UserInfo VALUES (1004,"Chan",{ })

```

```
INSERT INTO UserInfo VALUES (  
    1005,  
    "John",  
    {  
        "age":60,  
        "address":{  
            "street":"Taylor Blvd",  
            "city":"San Fransisco",  
            "state":"CA",  
            "phones":{"area":408,"number":50,"kind":"work"}  
        },  
        "expenses":{"housing":1000,"travel":300},  
        "connections":[30,5,null]  
    }  
)  
  
INSERT INTO UserInfo VALUES (  
    1006,  
    "Cathy",  
    {  
        "address":{  
            "street":"26th Avenue",  
            "city":"Chennai",  
            "state":"TN"  
        },  
        "vehicles":[  
            {  
                "vid":98642,  
                "vtype":"bike"  
            }  
        ]  
    }  
)
```

Query Optimization

This chapter discusses about query optimization in Oracle NoSQL Database.

This chapter contains the following topics:

- [Using Indexes for Query Optimization](#)
- [Finding Applicable Indexes](#)
- [Examples: Using Indexes for Query Optimization](#)
- [Choosing the Best Applicable Index](#)

Using Indexes for Query Optimization

In Oracle NoSQL Database, the query processor can identify which of the available indexes are beneficial for a query and rewrite the query to make use of such an index. "Using" an index means scanning a contiguous subrange of its entries, potentially applying further filtering conditions on the entries within this subrange, and using the primary keys stored in the surviving index entries to extract and return the associated table rows. The subrange of the index entries to scan is determined by the conditions appearing in the WHERE clause, some of which may be converted to search conditions for the index. Given that only a (hopefully small) subset of the index entries will satisfy the search conditions, the query can be evaluated without accessing each individual table row, thus saving a potentially large number of disk accesses.

Notice that in Oracle NoSQL Database, a primary-key index is always created by default. This index maps the primary key columns of a table to the physical location of the table rows. Furthermore, if no other index is available, the primary index will be used. In other words, there is no pure "table scan" mechanism; a table scan is equivalent to a scan via the primary-key index.

When it comes to indexes and queries, the query processor must answer two questions:

1. Is an index applicable to a query? That is, will accessing the table via this index be more efficient than doing a full table scan (via the primary index).
2. Among the applicable indexes, which index or combination of indexes is the best to use?

Regarding question (1), for queries with NESTED TABLES, secondary indexes will be considered for the target table only; in the current implementation, ancestor and/or descendant tables will always be accessed via their primary index.

Regarding question (2), the current implementation does not support index anding or index oring. As a result, the query processor will always use exactly one index (which may be the primary-key index). Furthermore, there are no statistics on the number and distribution of values in a table column or nested fields. As a result, the query processor has to rely on some simple heuristics in choosing among the applicable indexes. In addition, SQL for Oracle NoSQL Database allows for the inclusion of index hints in the queries, which are used as user instructions to the query processor about which index to use.

Finding Applicable Indexes

To find applicable indexes, the query processor looks at the conditions in the WHERE clause, trying to "match" such predicates with the index paths that define each index and convert the matched predicates to index search conditions. In general the WHERE clause consists of one or more conditions connected with AND or OR operators, forming a tree whose leaves are the conditions and whose internal nodes are the AND/OR operators. Let a predicate be any subtree of this WHERE-clause tree. The query processor will consider only top-level AND predicates, i.e., predicates that appear as the operands of a root AND node. If the WHERE clause does not have an AND root, the whole WHERE expression is considered a single top-level AND predicate. Notice that the query processor does not currently attempt to reorder the AND/OR tree in order to put it in conjunctive normal form. On the other hand, it does flatten the AND/OR tree so that an AND node will not have another AND node as a child, and an OR node will not have another OR node as a child. For example, the expression `a = 10 and b < 5 and (c > 10 or c < 0)` has 3 top-level AND predicates: `a = 10`, `b < 5`, and `(c > 10 or c < 0)`, whereas the expression `a = 10 and b < 5 and c > 10 or c < 0` has an OR as its root and the whole of it is considered as a single top-level AND predicate. For brevity, in the rest of this section we will use the term "predicate" to mean top-level AND predicate.

The query processor will also look at the expressions in the ORDER BY and GROUP BY clauses in order to find sorting indexes, that is, indexes that sort the table rows according to the expressions appearing in these clauses. As explained in sections [GROUP BY Clause](#) and [ORDER BY Clause](#), use of a sorting index will result in more efficient and memory-sparing sorting and grouping.

The query processor will consider an index applicable to a query if the index is a sorting one or if the query contains at least one index predicate: a predicate that can be evaluated during an index scan, using the content of the current index entry only, without the need to access the associated table row. Index predicates are further categorized as start/stop predicates or filtering predicates. A start/stop predicate participates in the establishment of the first/last index entry to be scanned during an index scan. A filtering predicate is applied during the index scan on the entries being scanned. In the current implementation, the following kinds of predicates are considered as candidate start/stop predicates:

- comparisons, using either the value or sequence (any) comparison operators, but not `!=` or `!=any`,
- IS NULL and IS NOT NULL operators,
- EXISTS and NOT EXISTS predicates, and
- IN predicates

However, if an index is created with the WITH NO NULLS clause, IS NULL and NOT EXISTS predicates cannot be used as index predicates for that index. In fact, such an index can be used by a query only if the query has an index predicate for each of the indexed fields.

An index is called a covering index with respect to a query if the query can be evaluated using only the entries of that index, that is, without the need to retrieve the associated rows.

If an index is used in a query, its index predicates are removed from the query because they are evaluated by the index scan. We say that index predicates are "pushed to the index". In the rest of this section we explain applicable indexes further via a number of example queries, and using the non-json indexes from the [Appendix](#). The algorithm for finding applicable json indexes is essentially the same as for non-json indexes. The same is true for geometry indexes, with the exception that geosearch predicates that are pushed to the index are not removed from the query, because they need to stay there to eliminate false positive results from the index scans.

Examples: Using Indexes for Query Optimization

Example 9-1 Using Indexes for Query Optimization

```
SELECT * FROM Users2
WHERE 10 < income AND income < 20
```

The query contains 2 index predicates. Indexes idx1, idx2, midx2, and midx3 are all applicable. For index idx1, 10 < income is a start predicate and income < 20 is a stop predicate. For the other indexes, both predicates are filtering predicates. If, say, idx2 were to be used, the subrange to scan is the whole index. Obviously, idx1 is better than the other indexes in this case. Notice however, that the number of table rows retrieved would be the same whether idx1 or idx2 were used. If midx2 or midx3 were used, the number of distinct rows retrieved would be the same as for idx1 and idx2, but a row would be retrieved as many times as the number of elements in the phones array of that row. Such duplicates are eliminated from the final query result set.

Notice that if index idx2 was created WITH NO NULLS, it would not be applicable to this query, because it does not have index predicates for fields address.state and address.city. For example, if Users2 contains a row where address.city is NULL and income is 15, the index would not contain any entry for this row, and as a result, if the index was used, the row would not appear in the result, even though it does qualify. The same is true for indexes midx2 and midx3. On the other hand, even if idx1 was created WITH NO NULLS, it would still be applicable, because it indexes a single field (income) and the query contains 2 start/stop predicates on that field.

Example 9-2 Using Indexes for Query Optimization

```
SELECT * FROM Users2
WHERE 20 < income OR income < 10
```

The query contains 1 index predicate, which is the whole WHERE expression. Indexes idx1, idx2, midx2, midx3 are all applicable. For all of them, the predicate is a filtering predicate.

Example 9-3 Using Indexes for Query Optimization

```
SELECT * FROM Users2
WHERE 20 < income OR age > 70
```

There is no index predicate in this case, because no index has information about user ages.

Example 9-4 Using Indexes for Query Optimization

```
SELECT * FROM Users2 u
WHERE u.address.state = "CA"
      AND u.address.city = "San Jose"
```

Only idx2 is applicable. There are 2 index predicates, both of which serve as both start and stop predicates.

Example 9-5 Using Indexes for Query Optimization

```
SELECT id, 2*income FROM Users2 u
WHERE u.address.state = "CA"
      AND u.address.city = "San Jose"
```

Only idx2 is applicable. There are 2 index predicates, both of which serve as both start and stop predicates. In this case, the id and income information needed in the SELECT clause is available in the index. As a result, the whole query can be answered from the index only, with no access to the table. We say that index idx2 is a covering index for the query in Example 5. The query processor will apply this optimization.

Example 9-6 Using Indexes for Query Optimization

```
SELECT * FROM Users2 u
WHERE u.address.state = "CA"
      AND u.address.city = "San Jose"
      AND u.income > 10
```

idx1, idx2, midx2, and midx3 are applicable. For idx2, there are 3 index predicates: the state and city predicates serve as both start and stop predicates; the income predicate is a start predicate. For idx1 only the income predicate is applicable, as a start predicate. For midx2 and midx3, the income predicate is a filtering one.

Example 9-7 Using Indexes for Query Optimization

```
SELECT * FROM Users2 u
WHERE u.address.state = "CA"
      AND u.income > 10
```

idx1, idx2, midx2, and midx3 are applicable. For idx2, there are 2 index predicates: the state predicate serves as both start and stop predicate; the income predicate is a filtering predicate. The income predicate is a start predicate for idx1 and a filtering predicate for midx2 and midx3.

Example 9-8 Using Indexes for Query Optimization

```
DECLARE $city STRING
```

```
SELECT * FROM Users2 u
WHERE u.address.state = "CA"
      AND u.address.city = $city
      AND (u.income > 50
           OR (10 < u.income
              AND u.income < 20))
```

idx1, idx2, midx2, and midx3 are applicable. For idx2, there are 3 index predicates. The state and city predicates serve as both start and stop predicates. The composite income predicate is a filtering predicate for all the applicable indexes (it's rooted at an OR node).

Example 9-9 Using Indexes for Query Optimization

```
SELECT u.address.city, SUM(u.expenses.values())
FROM Users2 u
```

```
WHERE u.address.state = "CA"
GROUP BY u.address.city
ORDER BY SUM(u.expenses.values())
```

In this example, for each city in California, the total amount of user expenditures in that city is returned. The query orders the results by the total amount. Only `idx2` is applicable. The state predicate is both a stop and a start predicate. Furthermore, the index is a sorting index, because for any given state it sorts the table rows by the names of the cities in that state and the `GROUP BY` groups by the cities in CA. As a result, the grouping in this query will be index-based and the `ORDER BY` will be generic. Notice that if instead of `idx2` there were 2 separate indexes, one on states and another on cities, both would be applicable: the first because of the state predicate, and the second because of the grouping. In this case, the query processor would choose the state index in order to reduce the number of rows accessed, at the expense of doing a generic `GROUP BY`.

Example 9-10 Using Indexes for Query Optimization

```
SELECT id FROM Users3 u
WHERE EXISTS u.info.income
```

In this example we use table `Users3`, which stores all information about users as json data. The query looks for users who record their income. Index `jidx1` is applicable. The `EXISTS` condition is actually converted to 2 index start/stop conditions: `u.info.income < EMPTY` and `u.info.income > EMPTY`. As a result, two range scans are performed on the index.

Example 9-11 Using Indexes for Query Optimization

```
SELECT * FROM users2 u
WHERE (u.address.state, u.address.city) IN
      (("CA", "San Jose"), ("MA", "Boston"))
```

In this example, the `idx2` index will be used. Two scans will be performed on the index: one for entries whose state and city fields are "CA" and "San Jose", respectively, and another for entries whose state and city fields are "MA" and "Boston", respectively.

Example 9-12 Using Indexes for Query Optimization

```
SELECT * FROM users2 u
WHERE u.address.state in ("CA", "MA") AND
      u.address.city in ("San Jose", "Boston")
```

In this example, the `idx2` index will be used. Four scans will be performed on the index. The search keys for these scans are determined by the cartesian product of the keys in the right-hand-side of the two `IN` operators: ("CA", "San Jose"), ("CA", "Boston"), ("MA", "San Jose"), and ("MA", "Boston").

Example 9-13 Using Indexes for Query Optimization

```
SELECT * FROM users2 u
WHERE (u.address.state, u.income) IN
      (("CA", 10000), ("MA", 20000))
```

In this example, the `idx2` index will be used. Two scans will be performed on the index: one for entries whose state field is "CA", and another for entries whose state field is "MA".

Furthermore, the whole IN condition will be used as a filtering predicate on the entries returned by the two scans.

As the above examples indicate, a predicate will be used as a start/stop predicate for an index IDX only if:

- It is of the form <path expr> op <const expr>, or <const expr> op <path expr>, or (<path expr1>, ... <path exprN>) IN (<const exprs>)
- op is a comparison operator (EXISTS, NOT EXISTS, IS NULL and IS NOT NULL are converted to predicates of this form, as shown in Q9).
- <const expr> is an expression built from literals and external variables only (does not reference any tables or internal variables)
- <path expr> is a path expression that is "matches" an index path P appearing in the CREATE INDEX statement for IDX. So far we have seen examples of exact matches only. In the examples below we will see some non-exact matches as well.
- If P is not IDX's 1st index path, there are equality start/stop predicates for each index path appearing before P in IDX's definition.
- The comparison operator may be one of the "any" operators. Such operators are matched against the multi-key index paths of multi-key indexes. As shown in the examples below, additional restrictions apply for such predicates.

Example 9-14 Using Indexes for Query Optimization

```
SELECT * FROM users2 u
WHERE u.connections[] = any 10
```

mdx1 is applicable and the predicate is both a start and a stop predicate.

Example 9-15 Using Indexes for Query Optimization

```
SELECT * FROM users2 u
WHERE u.connections[0:4] = any 10
```

mdx1 is applicable. The predicate to push down to mdx1 is u.connections[] =any 10, in order to eliminate users who are not connected at all with user 10. However, the original predicate (u.connections[0:4] =any 10) must be retained in the query to eliminate users who do have a connection with user 10, but not among their 5 strongest connections. This is an example where the query path expression does not match exactly the corresponding index path.

Example 9-16 Using Indexes for Query Optimization

```
SELECT * FROM users2 u
WHERE u.connections[] > any 10
```

mdx1 is applicable and the predicate is a start predicate.

Example 9-17 Using Indexes for Query Optimization

```
SELECT id FROM users2 u
WHERE 10 < any u.connections[]
      AND u.connections[] < any 100
```


midx1 is applicable, but although each predicate by itself is an index predicate, only one of them can actually be used as such. To see why, first notice that the query asks for users that have a connection with id greater than 10 and another connection (which may or may not be the same as the 1st one) with id less than 100. Next, consider a Users2 table with only 2 users (say with ids 200 and 500) having the following connections arrays respectively: [1, 3, 110, 120] and [1, 50, 130]. Both of these arrays satisfy the predicates in the query, and both users should be returned as a result. Now, consider midx1; it contains the following 7 entries:

[1, 200], [1, 500], [3, 200], [50, 500], [110, 200], [120, 200], [130, 500]

By using only the 1st predicate as a start predicate to scan the index, and applying the 2nd predicate on the rows returned by the index scan, the result of the query is 500, 200, which is correct. If on the other hand both predicates were used for the index scan, only entry [50, 500] would qualify, and the query would return only user 500.

Example 9-18 Using Indexes for Query Optimization

To search for users who have a connection in the range between 10 and 100, the following query can be used:

```
SELECT id FROM users2 u
WHERE exist u.connections
      [10 < $element AND $element < 100]
```

Assuming the same 2 users as in Example 13, the result of this query is user 500 only and both predicates can be used as index predicates (start and stop), because both predicates apply to the same array element. The query processor will indeed push both predicates to midx1.

Example 9-19 Using Indexes for Query Optimization

```
SELECT * FROM Users2 u
WHERE u.address.phones.area = any 650
      AND u.address.phones.kind = any "work"
      AND u.income > 10
```

This query looks for users whose income is greater than 10, and have a phone number with area code 650, and also have a work phone number (whose area code may not be 650). Index midx3 is applicable, but the address.phones.kind predicate cannot be used as an index predicate (for the same reason as in Example 13). Only the area code predicate can be used as a start/stop predicate and the income predicate as a filtering one. Indexes idx1, idx2, and midx2 are also applicable in Example 15.

Example 9-20 Using Indexes for Query Optimization

```
SELECT * FROM Users2 u
WHERE u.expenses.housing = 10000
```

idx4 is applicable and the predicate is both a start and a stop predicate. midx4 is also applicable. To use midx4, two predicates must be pushed to it, even though only one appears in the query. The 1st predicate is on the "keys" index field and the second on the "values" field. Specifically, the predicates key = "price" and value = 10000 are pushed as start/stop predicates. This is another example where the match between the query path expression and an index path is not exact: we match expenses.housing with the expenses.values() index path, and additionally, generate an index predicate for the properties.keys() index path.

Example 9-21 Using Indexes for Query Optimization

```
SELECT * FROM Users2 u
WHERE u.expenses.travel = 1000
      AND u.expenses.clothes > 500
```

midx4 is applicable. Each of the query predicates is by itself an index predicate and can be pushed to midx4 the same way as the expenses.housing predicate in the previous example. However, the query predicates cannot be both pushed (at least not in the current implementation). The query processor has to choose one of them to push and the other will remain in the query. Because the expenses.travel predicate is an equality one, it's more selective than the greater-than predicate and the query processor will use that.

Optimizing unnesting queries with the UNNEST clause

An UNNEST clause is recommended when there is an index on the array(s) or map(s) that are being unnested. The UNNEST clause places some restrictions on the kinds of expressions that it contains. These restrictions can help the query processor in choosing the appropriate index on the unnested arrays/maps, resulting in much better performance.

Note

See [Limitation for expression usage in the UNNEST clause](#) for more details on the restrictions placed by the UNNEST clause. For SQL statements to create tables and load data for using unnest queries, see [Example 5-3](#).

Example 9-22 Fetch different shows aired in the US and the number of people watching them

1. The query is specific to a country (US) and a specific show. So create an index on the country and showid fields of the stream_acct table.

```
CREATE INDEX idx_country_showid ON stream_acct( acct_data.country AS
string,

acct_data.contentStreamed[].showId AS integer)
WITH UNIQUE KEYS PER ROW
```

Note

The index must be created with the “unique keys per row property” in order for such indexes to be usable by queries that unnest the same array(s)/map(s) as the index.

2. The query to fetch different shows aired in the US and number of people watching it with the UNNEST clause.

```
SELECT $show.showId, count(*) as cnt
FROM stream_acct $s, unnest($s.acct_data.contentStreamed[] as $show)
WHERE $s.acct_data.country = "USA"
GROUP BY $show.showId ORDER BY count(*) desc
```

Output:

```
{"showId":15,"cnt":2}  
{"showId":16,"cnt":2}
```

The above query will use the `idx_country_showid` index. The country condition will be pushed to the index, the group-by will be index-based, and the index is a covering one for this query. An index that contains all required information to resolve the query is known as a **Covering Index** – it completely covers the query. Covering Index includes all the columns, the query refers to in the SELECT, JOIN, and WHERE clauses. If the UNNEST clause is not used, the index will not be considered.

Example 9-23 For every show aired by the application, the total watch time for all users

The following query returns, for each show, the total time users have spent watching the show.

```
SELECT $show.showId, sum($show.seriesInfo.episodes.minWatched)  
AS total_time FROM stream_acct $s,  
unnest($s.acct_data.contentStreamed[] AS $show)  
GROUP BY $show.showId  
ORDER BY sum($show.seriesInfo.episodes.minWatched)
```

Let us examine the effect of creating an index on `acct_data.contentStreamed[].showId` as the data is been grouped based on `showId`.

Despite the use of the UNNEST clause, this query cannot use the `idx_showid` index. This is because of the argument to the `sum()` function. The `idx_showid` index contains just the `showId` (and the primary key). So, the expression `$show.seriesInfo.episodes.minWatched` cannot be evaluated from the index. There are two ways to optimize this query.

Option 1: Create an additional index:

Create an additional composite index on the `showId` and `minWatched` fields as both are used in the query.

```
CREATE INDEX idx_showid_minWatched ON  
stream_acct(acct_data.contentStreamed[].showId AS integer,  
acct_data.contentStreamed[].seriesInfo[].episodes[].minWatched AS  
integer,  
acct_data.contentStreamed[].seriesInfo[].episodes[].episodeID as  
integer)  
WITH UNIQUE KEYS PER ROW
```

Note

The `episodeID` must be added in this index, as the last index path, in order for the "unique keys per row" constraint to be satisfied. This index will be used by the query, as a covering index.

Option 2: Avoid the cost of an additional index:

You can rewrite the query to use the `idx_showid` index. The rewritten query below uses the `idx_showid` index, but the index is not covering.

```
SELECT $show.showId, sum($s.acct_data.contentStreamed[$element.showId
= $show.showId].
                        seriesInfo.episodes.minWatched) AS total_time
FROM stream_acct $s, unnest($s.acct_data.contentStreamed[] AS $show) GROUP
BY $show.showId
ORDER BY sum($s.acct_data.contentStreamed[$element.showId
= $show.showId].seriesInfo.episodes.minWatched)
```

Output:

```
{"showId":26,"total_time":225}
{"showId":16,"total_time":440}
{"showId":15,"total_time":642}
```

Example 9-24 The total watch time of users per show and season

```
SELECT $show.showId, $seriesInfo.seasonNum,
sum($seriesInfo.episodes.minWatched) AS length
FROM stream_acct n,
unnest(n.acct_data.contentStreamed[] AS $show, $show.seriesInfo[]
as $seriesInfo)
GROUP BY $show.showId, $seriesInfo.seasonNum
ORDER BY sum($seriesInfo.episodes.minWatched)
```

Output:

```
{"showId":26,"seasonNum":2,"length":80}
{"showId":26,"seasonNum":1,"length":145}
{"showId":16,"seasonNum":2,"length":190}
{"showId":16,"seasonNum":1,"length":250}
{"showId":15,"seasonNum":2,"length":295}
{"showId":15,"seasonNum":1,"length":347}
```

For best performance of the above query, create the following index. The index `idx_showid_seasonNum_minWatched` is a multi key index. The `episodeID` must be added in this index, as the last index path, in order for the "unique keys per row" constraint to be satisfied.

```
CREATE INDEX idx_showid_seasonNum_minWatched ON
stream_acct(acct_data.contentStreamed[].showId as integer,
            acct_data.contentStreamed[].seriesInfo[].seasonNum as integer,
            acct_data.contentStreamed[].seriesInfo[].episodes[].minWatched as
integer,
            acct_data.contentStreamed[].seriesInfo[].episodes[].episodeID as
integer)
WITH UNIQUE KEYS PER ROW
```

Note

An index is called a multikey index if for each row of data in the table, there are multiple entries created in the index. In a multikey index, there is at least one index path that uses `.keys()`, `.values()`, or `[]` steps. Any such index path will be called a multikey index path.

If you want your query to use an index on the `array(s)/maps()` that it is unnesting, each path expression in the `UNNEST` clause must match with the `multikey_path_prefix` of an index path in the index. As shown in the syntax for the [CREATE INDEX Statement](#) the `multikey_path_prefix` is the part of an index path up to and including the last multikey step. For example, the `multikey_path_prefix` of the first index path should match the first path expression in the `UNNEST` clause and so on. If this is not the case, the index will not be used.

In the above query the expression `n.value.contentStreamed[]` matches the `multikey_path_prefix` of the first index path in `idx_showid_seasonNum_minWatched`, and the expression `$show.seriesInfo[]` matches the `multikey_path_prefix` of the second index path, after the `$show` variable is replaced with its domain expression. So this index will be used in the query.

Choosing the Best Applicable Index

To choose an index for a query, the query processor uses a simple heuristic together with any user-provided index hints.

Syntax

```
hints ::= '/*+' hint* '*/'
```

```
hint ::= (
    (PREFER_INDEXES "(" name_path index_name* ")") |
    (FORCE_INDEX "(" name_path index_name ")") |
    (PREFER_PRIMARY_INDEX "(" name_path ")") |
    (FORCE_PRIMARY_INDEX "(" name_path ")")
) [STRING]
```

There are 2 kinds of hints: a `FORCE_INDEX` hint and a `PREFER_INDEXES` hint. The `FORCE_INDEX` hint specifies a single index and the query is going to use that index without considering any of the other indexes (even if there are no index predicates for the forced index). The `PREFER_INDEXES` hint specifies one or more indexes. The query processor may or may not use one of the preferred indexes. Specifically, in the absence of a forced index, index selection works as follows.

The query processor uses the heuristic to assign a score to each applicable index and then chooses the one with the highest score. If two or more indexes have the same score, the index chosen is the one whose name is alphabetically before the others. In general, preferred indexes will get high scores, but it is possible that other indexes may still win. Describing the details of the heuristic is beyond the scope of this document, but a few high-level decisions are worth mentioning:

- If the query has a complete primary key, the primary index is used.
- Indexes that are preferred (via a `PREFER` hint), covering, or have a complete key (i.e., there is an equality predicate on each of its index fields) get high scores and will normally prevail over other indexes.

- Among 2 indexes where one is a sorting index, the other is not, and the 2 indexes would otherwise have the same score, the sorting index is chosen.

The `FORCE_INDEX` and `PREFER_INDEXES` hints specify indexes by their name. Since the primary index has no explicit name, 2 more hints are available to force or to prefer the primary index: `FORCE_PRIMARY_INDEX` and `PREFER_PRIMARY_INDEX`. Hints are inserted in the query as a special kind of comment that appears immediately after the `SELECT` keyword. Here is the relevant syntax:

The '+' character immediately after (with no spaces) the comment opening sequence ('/*') is what turns the comment into a hint. The string at the end of the hint is just for informational purposes (a comment for the hint) and does not play any role in the query execution.

Appendix

The following code creates `Users2` and `User3` tables.

```
CREATE TABLE Users2 (
  id INTEGER,
  income INTEGER,
  address RECORD(
    street STRING,
    city STRING,
    state STRING,
    phones ARRAY(
      RECORD(
        area INTEGER,
        number INTEGER,
        kind STRING
      )
    )
  ),
  connections ARRAY(INTEGER),
  expenses MAP(INTEGER),
  PRIMARY KEY (id)
)
```

```
CREATE TABLE users3 (id INTEGER, info JSON, PRIMARY KEY(id))
```

The following code populates `Users2` and `User3` tables with sample rows.

```
INSERT INTO Users2 VALUES (
  0,
  1000,
  {
    "street" : "somewhere",
    "city": "Boston",
    "state" : "MA",
    "phones" : [
      { "area":408, "number":50, "kind":"work" },
      { "area":415, "number":60, "kind":"work" },
      { "area":NULL, "number":52, "kind":"home" }
    ]
  },
  [],
  {},
  {}
)
```

```

        [ 100, 20, 20, 10, 20],
        { "housing" : 1000, "clothes" : 230, "books" : 20 }
    )

INSERT INTO Users2 VALUES (
    1,
    NULL,
    {
        "street" : "everywhere",
        "city": "San Fransisco",
        "state" : "CA",
        "phones" : [
            { "area":408, "number":50, "kind":"work" },
            { "area":408, "number":60, "kind":"home" }
        ]
    },
    [],
    { "housing" : 1000, "travel" : 300 }
)

INSERT INTO Users2 VALUES (
    2,
    2000,
    {
        "street" : "nowhere",
        "city": "San Jose",
        "state" : "CA",
        "phones" : [ ]
    },
    NULL,
    NULL
)

INSERT INTO users3 VALUES (
    0,
    {
        "income" : 1000,
        "address": {
            "street" : "somewhere",
            "city": "Boston",
            "state" : "MA",
            "phones" : [
                { "area":408, "number":50, "kind":"work" },
                { "area":415, "number":60, "kind":"work" },
                { "area":null, "number":52, "kind":"home" }
            ]
        },
        "expenses" : { "housing" : 1000, "clothes" : 230, "books" : 20 },
        "connections" : [ 100, 20, 20, 10, 20]
    }
)

INSERT INTO users3 VALUES (
    1,
    {

```

```

        "income" : null,
        "address": {
            "street" : "everywhere",
            "city": "San Fransisco",
            "state" : "CA",
            "phones" : [
                { "area":408, "number":50, "kind":"work" },
                { "area":408, "number":60, "kind":"home" },
                "4083451232"
            ]
        },
        "expenses" : { "housing" : 1000, "travel" : 300 },
        "connections" : [ ]
    }
)

INSERT INTO users3 VALUES (
    2,
    {
        "income" : 2000,
        "address": {
            "street" : "nowhere",
            "city": "San Jose",
            "state" : "CA",
            "phones" : [ ]
        },
        "expenses" : null,
        "connections" : null
    }
)

INSERT INTO users3 VALUES (3,{})

INSERT INTO users3 VALUES (
    4,
    {
        "address": {
            "street" : "top of the hill",
            "city": "San Fransisco",
            "state" : "CA",
            "phones" : { "area":408, "number":50, "kind":"work" }
        },
        "expenses" : { "housing" : 1000, "travel" : 300},
        "connections" : [ 30, 5, null ]
    }
)

INSERT INTO Users3 VALUES (
    5,
    {
        "address": {
            "street" : "end of the road",
            "city": "Portland",
            "state" : "OR"
        }
    }
)

```



```
    }  
  )
```

The following are some examples of indexes.

Example 9-25 Simple Index

```
CREATE INDEX idx1 ON Users2 (income)
```

It creates an index with one entry per user in the Users table. The entry contains the income and id of the user represented by the row. The contents of this index for the sample rows in Users2 are:

```
[ 1000, 0 ]  
[ 2000, 2 ]  
[ NULL, 1 ]
```

If the WITH NO NULLS clause were used in the above create index statement, the last of the above 3 entries would not appear in the index.

Example 9-26 Simple Index

```
CREATE INDEX idx2 ON Users2 (address.state, address.city, income)
```

It creates an index with one entry per user in the Users table. The entry contains the state, city, income and id of the user represented by the row. The contents of this index for the sample rows in Users2 are:

```
[ "CA", "San Fransisco", NULL, 1 ]  
[ "CA", "San Jose", 2000, 2 ]  
[ "MA", "Boston", 1000, 0 ]
```

Example 9-27 Simple Index

```
CREATE INDEX idx3 ON Users2 (expenses.books)
```

Creates an index entry for each user. The entry contains the user's spending on books, if the user does record spending on books, or EMPTY if there is no "books" entry in expenses, or NULL if there is no expenses map at all (i.e. the value of the expenses column is NULL). The contents of this index for the sample rows in Users2 are:

```
[ 20, 0 ]  
[ EMPTY, 1 ]  
[ NULL, 2 ]
```

If the WITH NO NULLS clause were used in the above create index statement, only the first of the above 3 entries would appear in the index.

Example 9-28 Simple Index

```
CREATE INDEX idx4 ON users2 (expenses.housing, expenses.travel)
```

Creates an index entry for each user. The entry contains the user's housing expenses, or EMPTY if the user does not record housing expenses, and the user's travel expenses, or EMPTY if the user does not record travel expenses. If expenses is NULL, both fields in the index entry will be NULL. The contents of this index for the sample rows in Users2 are:

```
[ 1000, 300, 1 ]  
[ 1000, EMPTY, 0 ]  
[ NULL, NULL, 2 ]
```

Example 9-29 Multi-Key Index

```
CREATE INDEX midx1 ON Users2 (connections[])
```

Creates an index on the elements of the connections array. The contents of this index for the sample rows in Users2 are:

```
[ 10, 0 ]  
[ 20, 0 ]  
[ 100, 0 ]  
[ EMPTY, 1 ]  
[ NULL, 2 ]
```

If the WITH NO NULLS clause were used in the above create index statement, the last 2 of the above entries would not appear in the index.

Example 9-30 Multi-Key Index

```
CREATE INDEX midx2 ON Users2 (address.phones[].area, income)
```

Creates an index on the area codes and income of users. The contents of this index for the sample rows in Users2 are:

```
[ 408, 1000, 0 ]  
[ 408, NULL, 1 ]  
[ 415, 1000, 0 ]  
[ EMPTY, 2000, 2 ]  
[ NULL, 1000, 0 ]
```

Example 9-31 Multi-Key Index

```
CREATE INDEX midx3 ON Users2  
  (address.phones[].area, address.phones[].kind, income)
```

Creates an index on the area codes, the phone number kinds, and the income of users. The contents of this index for the sample rows in Users2 are:

```
[ 408, "work", 1000, 0 ]  
[ 408, "home", NULL, 1 ]  
[ 408, "work", NULL, 1 ]  
[ 415, "work", 1000, 0 ]  
[ EMPTY, EMPTY, 2000, 2 ]  
[ NULL, "home", 1000, 0 ]
```

Example 9-32 Multi-Key Index

```
CREATE INDEX midx4 ON Users2 (  
    expenses.keys(), expenses.values())
```

Creates an index on the fields (both keys and values) of the expenses map. The contents of this index for the sample rows in Users2 are:

```
[ "books", 50, 0 ]  
[ "clothes", 230, 0 ]  
[ "housing", 1000, 0 ]  
[ "housing", 1000, 1 ]  
[ "travel", 300, 1 ]  
[ NULL, NULL, 2 ]
```

Example 9-33 Simple Typed json Index

```
CREATE INDEX jidx1 ON users3(info.income AS INTEGER)
```

It creates an index with one entry per user in the Users table. The entry contains the income and id (the primary key) of the user represented by the row. The contents of this index for the sample rows in Users3 are:

```
[ 1000, 0 ]  
[ 2000, 2 ]  
[ EMPTY, 4 ]  
[ EMPTY, 5 ]  
[ JNULL, 1 ]  
[ NULL, 3 ]
```

Example 9-34 Simple Typed json Index

```
CREATE INDEX jidx1u ON users3 (  
    info.income AS ANYATOMIC)
```

It creates an untyped index on info.income. The contents of this index are the same as in jidx1 above, but the values 1000 and 200 are stored as Numbers instead of integers. If the following row is added to the users3 table:

```
INSERT INTO users3 VALUES (  
    6,  
    {  
        "address": {},  
        "expenses" : {},  
        "connections" : []  
    }  
)
```

The index will look like this:

```
[ "none", 6 ]
```

```
[ EMPTY, 5 ]
[ EMPTY, 4 ]
[ NULL, 3 ]
[ 2000, 2 ]
[ JNULL, 1 ]
[ 1000, 0 ]
```

Example 9-35 Simple Typed json Index

```
CREATE INDEX jidx2 ON users3 (
    info.address.state AS STRING,
    info.address.city AS STRING,
    info.income AS INTEGER)
```

It creates an index with one entry per user in the Users table. The entry contains the state, city, income and id (the primary key) of the user represented by the row. The contents of this index for the sample rows in Users3 are:

```
[ "CA", "San Fransisco", EMPTY, 4 ]
[ "CA", "San Fransisco", JNULL, 1 ]
[ "CA", "San Jose", 2000, 2 ]
[ "MA", "Boston", 1000, 0 ]
[ "OR", "Portland", EMPTY, 5 ]
[ NULL, NULL, NULL, 3 ]
```

Example 9-36 Simple Typed json Index

```
CREATE INDEX jidx3 ON users3 (
    info.expenses.books AS INTEGER)
```

Creates an index entry for each user. The entry contains the user's spending on books, if the user does record spending on books, or EMPTY if there is no "books" entry in expenses or there is no expenses map at all, or NULL if there is no info at all (i.e. the value of the info column is NULL). The contents of this index for the sample rows in Users3 are:

```
[ 20, 0 ]
[ EMPTY, 1 ]
[ EMPTY, 2 ]
[ EMPTY, 4 ]
[ EMPTY, 5 ]
[ NULL, 3 ]
```

Example 9-37 Simple Typed json Index

```
CREATE INDEX jidx4 ON users3 (
    info.expenses.housing AS INTEGER,
    info.expenses.travel AS INTEGER)
```

Creates an index entry for each user. The entry contains 2 fields: (a) the user's housing expenses, or EMPTY if the user does not record housing expenses or there is no expenses field at all, and (b) the user's travel expenses, or EMPTY if the user does not record travel

expenses or there is no expenses field at all. If info is NULL, both fields in the index entry will be NULL. The contents of this index for the sample rows in Users3 are:

```
[ 1000, 300, 1 ]
[ 1000, 300, 4 ]
[ 1000, EMPTY, 0 ]
[ EMPTY, EMPTY, 2 ]
[ EMPTY, EMPTY, 5 ]
[ NULL, NULL, 3 ]
```

Example 9-38 Multi-Key Typed json Index

```
CREATE INDEX jmidx1 ON users3 (
    info.connections[] AS INTEGER)
```

Creates an index on the elements of the connections array. The contents of this index for the sample rows in Users3 are:

```
[ 5, 4 ]
[ 10, 0 ]
[ 20, 0 ]
[ 30, 4 ]
[ 100, 0 ]
[ EMPTY, 1 ]
[ EMPTY, 5 ]
[ JNULL, 2 ]
[ JNULL, 4 ]
[ NULL, 3 ]
```

Example 9-39 Multi-Key Typed json Index

```
CREATE INDEX jmidx2 ON users3 (
    info.address.phones[].area AS INTEGER,
    info.income AS INTEGER)
```

Creates an index on the area codes and income of users. The contents of this index for the sample rows in Users3 are:

```
[ 408, 1000, 0 ]
[ 408, EMPTY, 4 ]
[ 408, JNULL, 1 ]
[ 415, 1000, 0 ]
[ EMPTY, 2000, 2 ]
[ EMPTY, EMPTY, 5 ]
[ EMPTY, JNULL, 1 ]
[ JNULL, 1000, 0 ]
[ NULL, NULL, 3 ]
```

Example 9-40 Multi-Key Typed json Index

```
CREATE INDEX jmidx2u ON users3 (
    info.address.phones[].area AS ANYATOMIC,
    info.income AS INTEGER)
```

This is a variation of the jmidx2 index, where the first index path is untyped and second is typed. The contents of jmidx2 and jmidx2u are the same, except that in jmidx2u the numeric values in the first column are stored as Numbers instead of integers.

Example 9-41 Multi-Key Typed json Index

```
CREATE INDEX jmidx3 ON users3 (
    info.address.phones[].area AS INTEGER,
    info.address.phones[].kind AS string,
    info.income AS INTEGER)
```

Creates an index on the area codes, the phone number kinds, and the income of users. The contents of this index for the sample rows in Users3 are:

```
[ 408, "home", JNULL, 1 ]
[ 408, "work", 1000, 0 ]
[ 408, "work", EMPTY, 4 ]
[ 408, "work", JNULL, 1 ]
[ 415, "work", 1000, 0 ]
[ EMPTY, EMPTY, 2000, 2 ]
[ EMPTY, EMPTY, EMPTY, 5 ]
[ EMPTY, EMPTY, JNULL, 1 ]
[ JNULL, "home", 1000, 0 ]
[ NULL, NULL, NULL, 3 ]
```

Example 9-42 Multi-Key Typed json Index

```
CREATE INDEX jmidx4 ON users3 (
    info.expenses.keys(),
    info.expenses.values() AS INTEGER)
```

Creates an index on the fields (both keys and values) of the expenses map. Notice that the keys() portion of the index definition must not declare a type. This is because the type will always be String. The contents of this index for the sample rows in Users2 are:

```
[ "books", 50, 0 ]
[ "clothes", 230, 0 ]
[ "housing", 1000, 0 ]
[ "housing", 1000, 1 ]
[ "housing", 1000, 4 ]
[ "travel", 300, 1 ]
[ "housing", 1000, 4 ]
[ EMPTY, EMPTY, 2 ]
[ EMPTY, EMPTY, 5 ]
[ NULL, NULL, 3 ]
```


Query Plan

A query execution plan is the sequence of operations Oracle NoSQL Database performs to run a query.

Topics:

- [Overview of a query plan](#)
- [Examples of query execution plan](#)

Overview of a query plan

Internally, a query execution plan is structured as a tree of plan iterators.

Each kind of iterator evaluates a different kind of expression that may appear in a query. In general the choice of index and the kind of associated index predicates can have a drastic effect on the query performance. As a result, you as a developer often want to see what index is used by a query and what predicates have been pushed down to it. Based on this information, you may want to force the use of a different index via index hints. This information is contained in the query execution plan. All Oracle NoSQL drivers provide APIs to display the execution plan of a query. All Oracle NoSQL graphical UIs including the IntelliJ, VSCode, and Eclipse plugins along with the Oracle Cloud Infrastructure Console include controls for displaying the query execution plan.

Iterators in a query execution plan

Some of the most common and important iterators used in queries are :

TABLE iterator

A TABLE iterator is responsible for:

- Scanning the index used by the query (which may be the primary index).
- Applying any filtering predicates pushed to the index.
- Retrieving the rows pointed to by the qualifying index entries if necessary. If the index is covering, the result set of the TABLE iterator is a set of index entries, otherwise it is a set of table rows.

Note

An index is called a covering index with respect to a query if the query can be evaluated using only the entries of that index, that is, without the need to retrieve the associated rows.

A TABLE iterator will always have the following properties:

- **target table:** The name of the target table in the query.
- **index used:** The name of the index used by the query. If the primary index were used, “primary index” would appear as the value of this property.

- **covering index:** Whether the index is covering or not.
- **row variable:** The name of a variable ranging over the table rows produced by the TABLE iterator. If the index is covering, no table rows are produced and this variable is not used.
- **index scans:** Contains the start and stop conditions that define the index scans to be performed.

A TABLE iterator has 2 more optional properties:

- **index row variable:** The name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the index variable will be bound to that entry.
- **index filtering predicate:** A predicate evaluated on every index entry produced by the index scan. If the result of this evaluation is true, the index variable is bound to this entry and the entry or its associated table row is returned as the result of the next() call on the TABLE iterator. Otherwise, the entry is skipped, the next entry from the index scan is produced, the predicate is evaluated again on this entry and it continues until a qualifying entry is found.

SELECT iterator

It is responsible for executing the SELECT expression.

RECEIVE iterator

It is a special internal iterator that separates the query plan into 2 parts:

- The RECEIVE iterator itself and all iterators that are above it in the iterator tree are executed at the driver.
- All iterators below the RECEIVE iterator are executed at the replication nodes (RNs); these iterators form a subtree rooted at the unique child of the RECEIVE iterator.

In general, the RECEIVE iterator acts as a **query coordinator**. It sends its subplan to appropriate RNs for execution and collects the results. It may perform additional operations such as sorting and duplicate elimination, and propagates the results to its ancestor iterators (if any) for further processing.

Distribution kinds :

A distribution kind specifies how the query will be distributed for execution across the RNs participating in an Oracle NoSQL database (a store). The distribution kind is a property of the RECEIVE iterator.

Different choices of Distribution kinds are:

- **SINGLE_PARTITION:** A SINGLE_PARTITION query specifies a complete shard key in its WHERE clause. As a result, its full result set is contained in a single partition, and the RECEIVE iterator will send its subplan to a single RN that stores that partition. A SINGLE_PARTITION query may use either the primary-key index or a secondary index.
- **ALL_PARTITIONS:** Queries use the primary-key index here and they don't specify a complete shard key. As a result, if the store has M partitions, the RECEIVE iterator will send M copies of its subplan to be executed over one of the M partitions each. See show topology to determine the number of partitions in your store.
- **ALL_SHARDS:** Queries use a secondary index here and they don't specify a complete shard key. As a result, if the store has N shards, the RECEIVE iterator will send N copies of its subplan to be executed over one of the N shards each.

Anatomy of a query execution plan:

Query execution takes place in batches. When a query subplan is sent to a partition or shard for execution, it will execute there until a batch limit is reached. For an on-premises NoSQL database, batch limit is the number of local results produced from the underlying partition/shard. The default is 100 results, but you can change it via a query-level option. For NoSQL Database Cloud Service, the batch limit is the number of read units consumed locally by the query. The default is 2000 read units (about 2MB of data), and it can only be decreased via a query-level option.

When the batch limit is reached, any local results that were produced are sent back to the RECEIVE iterator for further processing along with a boolean flag that says whether more local results may be available. If the flag is true, the reply includes resume information. If the RECEIVE iterator decides to resend the query to the same partition/shard, it will include this resume information in its request, so that the query execution will restart at the point where it stopped during the previous batch. This is because no query state is maintained at the RN after a batch finishes. The next batch for the same partition/shard may take place at the same RN as the previous batch or at a different RN that also stores the same partition/shard.

Examples of query execution plan

You can write some queries using the users table and understand how query execution plan is generated.

Description of the users table:

```
CREATE TABLE users (  
    id INTEGER,  
    firstName STRING,  
    lastName STRING,  
    otherNames ARRAY(RECORD(first STRING, last STRING)),  
    age INTEGER,  
    income INTEGER,  
    address JSON,  
    connections ARRAY(INTEGER),  
    expenses MAP(INTEGER),  
    PRIMARY KEY (id)  
)
```

The following index has been created in the users table.

```
CREATE INDEX idx_state_city_income on Users(address.state as string,  
    address.city as string, income)
```

Some examples of query execution plan :

- [Example 1 : Using a covering index in a query plan with only index scans](#)
- [Example 2 : Using a covering index in a query plan with index scans and index predicates](#)
- [Example 3: Using a non-covering index in a query plan with index scans](#)
- [Example 4: Sort the data using a Covering index](#)
- [Example 5: Sort the data using a field not part of the index](#)
- [Example 6: Group the data using a Covering index](#)

- [Example 7: Group data with fields not part of the index](#)

Example 1 : Using a covering index in a query plan with only index scans

An index is called a covering index with respect to a query if the query can be evaluated using only the entries of that index, that is, without the need to retrieve the associated rows.

Example 10-1 Fetch the id and income of users whose state is CA and their city value must be greater or equal to "S" and whose income is between 1000 and 2000

```
SELECT id, income FROM Users u WHERE u.address.state = "CA" AND
u.address.city >= "S" AND 1000 < income and income < 2000
```

Query execution plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_SHARDS",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "users",
      "row variable" : "$$u",
      "index used" : "idx_state_city_income",
      "covering index" : true,
      "index row variable" : "$$u_idx",
      "index scans" : [
        {
          "equality conditions" :
          { "address.state": "CA", "address.city": "Santaclara" },
          "range conditions" : { "income" : { "start value" : 1000, "start
inclusive" : false, "end value" : 2000, "end inclusive" : false } }
        }
      ]
    },
    "FROM variable" : "$$u_idx",
    "SELECT expressions" : [
      {
        "field name" : "id",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "#id",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u_idx"
          }
        }
      },
      {
        "field name" : "income",
```

```

    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "income",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u_idx"
      }
    }
  }
]
}

```

Explanation of the query execution plan :

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator. The only property of the RECEIVE iterator in this example, is the distribution kind whose value is **ALL_SHARDS**.
- The index **idx_state_city_income** is used here and this is a covering index as all the fields in the SELECT expression can be fetched only using the index entries. .
- The index scan property contains the start and stop conditions that define the index scans to be performed.

```

"index scans" : [
  {
    "equality conditions" :
    {"address.state":"CA","address.city":"Santaclara"},
    "range conditions" : { "income" : { "start value" : 1000, "start
inclusive" : false, "end value" : 2000, "end inclusive" : false } }
  }
]

```

In this query, only one index scan will be performed. The equality conditions correspond to the predicates `u.address.state = "CA"` and `u.address.city = "Santaclara"` from the query. The range conditions correspond to the predicates `1000 < income` and `income < 2000`. The index scan will start at the first entry whose `address.state` field is equal to **CA**, its `address.city` field is equal to **Santaclara**, and its `income` field is greater than 1000. The index scan will return all subsequent entries until the first entry whose `address.state` field is not **CA**, or its `address.city` field is not **Santaclara**, or its `income` field is greater than or equal to 2000.

- The index row variable is **\$\$u_idx**, which is the name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the **\$\$u_idx** variable will be bound to that entry.
- When the FROM iterator is a TABLE iterator, the FROM variable is the same as either the **index row variable** or the **row variable** of the TABLE iterator, depending on whether the used index is covering or not. In this example, the FROM variable is same as the index row variable (**\$\$u_idx**) as the index is covering because all the fields in the SELECT expression can be evaluated only using the index entries.
- This index row variable (**\$\$u_idx**) will be referenced by iterators implementing the other clauses of the SELECT expression.

- In the SELECT expression, two fields (**id** and **income**) are fetched. These correspond to two field names and field expressions in the SELECT expression clause.

```
{
  "field name" : "id",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "#id",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u_idx"
    }
  }
}
```

For every field to be fetched by the SELECT expression, the field expression is computed by a FIELD_STEP iterator. The FIELD_STEP iterator extracts and returns the value of a field (id as shown above) from the records produced by its input iterator. The same is done for every field to be fetched in the SELECT expression.

Example 2 : Using a covering index in a query plan with index scans and index predicates

Example 10-2 Fetch the id and income of users whose state is CA and whose income is greater than 2000

```
SELECT id, income FROM Users u WHERE u.address.state = "CA" AND income > 2000
```

Query execution plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_SHARDS",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "users",
      "row variable" : "$$u",
      "index used" : "idx_state_city_income",
      "covering index" : true,
      "index row variable" : "$$u_idx",
      "index scans" : [
        {
          "equality conditions" : {"address.state":"CA"},
          "range conditions" : {}
        }
      ],
      "index filtering predicate" :

```

```

{
  "iterator kind" : "GREATER_THAN",
  "left operand" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "income",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u_idx"
    }
  },
  "right operand" :
  {
    "iterator kind" : "CONST",
    "value" : 2000
  }
}
},
"FROM variable" : "$$u_idx",
"SELECT expressions" : [
  {
    "field name" : "id",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "#id",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u_idx"
      }
    }
  },
  {
    "field name" : "income",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "income",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u_idx"
      }
    }
  }
]
}
}

```

Explanation of the query execution plan :

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator. The only property of the RECEIVE iterator in this example, is the distribution kind whose value is **ALL_SHARDS**.

- The index **idx_state_city_income** is used here and this is a covering index as all the fields in the SELECT expression can be fetched only using the index entries.
- The index scan property contains the start and stop conditions that define the index scans to be performed.

```
"index scans" : [
  {
    "equality conditions" : {"address.state":"CA"},
    "range conditions" : {}
  }
]
```

In this example, only one index scan will be performed. The conditions correspond to the predicates `u.address.state = "CA"` from the query. Specifically, the starting index entry must have the value **CA** on the `address.state` field. All subsequent entries must have **CA** as the value of their `address.state` field, and the scan will stop as soon as an entry with a different state value is encountered. Although the query contains a range predicate on income, this predicate does not appear as a range condition of the index scan. This is because there is no equality condition on the `address.city` field that appears before the income field in the index definition, and as a result, the income predicate cannot be used to determine the boundaries of the scan. Instead, the income predicate can be used as an index filtering predicate that is applied on every index entry produced by the index scan.

- The index filtering predicate evaluates the filter criteria on the **income** field. Using the greater than operator the filter condition is evaluated.

```
"index filtering predicate" :
{
  "iterator kind" : "GREATER_THAN",
  "left operand" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "income",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u_idx"
    }
  },
  "right operand" :
  {
    "iterator kind" : "CONST",
    "value" : 2000
  }
}
```

- The index row variable is **\$\$u_idx** which is the name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the `$$u_idx` variable will be bound to that entry.
- When the FROM iterator is a TABLE iterator, the FROM variable is the same as either the **index row variable** or the **row variable** of the TABLE iterator, depending on whether the used index is covering or not. In this example, the FROM variable is same as the index row

variable (`$$u_idx`) as the index is covering because all the fields in the `SELECT` expression can be evaluated only using the index entries.

- This index row variable (`$$u_idx`) will be referenced by iterators implementing the other clauses of the `SELECT` expression.
- In the `SELECT` expression, two fields (`id` and `income`) are fetched. These correspond to two field names and field expressions in the `SELECT` expression clause.

```
{
  "field name" : "id",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "#id",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u_idx"
    }
  }
}
```

For every field to be fetched by the `SELECT` expression, the field expression is computed by a `FIELD_STEP` iterator. The `FIELD_STEP` iterator extracts and returns the value of a field (`id` as shown above) from the records produced by its input iterator. The same is done for every field to be fetched in the `SELECT` expression.

Example 3: Using a non-covering index in a query plan with index scans

An index becomes non-covering when query cannot be fully evaluated using only the entries of an index.

Example 10-3 Fetch id, age, and income of users residing in CA whose income is greater than 5000

```
SELECT id, age, income FROM Users u WHERE u.address.state = "CA"
AND income >5000
```

Query execution plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_SHARDS",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "users",
      "row variable" : "$$u",
      "index used" : "idx_state_city_income",
      "covering index" : false,
      "index row variable" : "$$u_idx",
```



```

"index scans" : [
  {
    "equality conditions" : {"address.state":"CA"},
    "range conditions" : {}
  }
],
"index filtering predicate" :
{
  "iterator kind" : "GREATER_THAN",
  "left operand" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "income",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u_idx"
    }
  },
  "right operand" :
  {
    "iterator kind" : "CONST",
    "value" : 5000
  }
}
},
"FROM variable" : "$$u",
"SELECT expressions" : [
  {
    "field name" : "id",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "id",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u"
      }
    }
  },
  {
    "field name" : "age",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "age",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u"
      }
    }
  },
  {
    "field name" : "income",

```

```

        "field expression" :
        {
            "iterator kind" : "FIELD_STEP",
            "field name" : "income",
            "input iterator" :
            {
                "iterator kind" : "VAR_REF",
                "variable" : "$$u"
            }
        }
    }
}
]
}

```

Explanation of the query execution plan :

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator. The only property of the RECEIVE iterator in this example, is the distribution kind whose value is **ALL_SHARDS**.
- The index **idx_state_city_income** is used here and in this example, it is not a covering index as the age field to be fetched is not part of the index entry.
- The index scan property contains the start and stop conditions that define the index scans to be performed.

```

"index scans" : [
    {
        "equality conditions" : {"address.state":"CA"},
        "range conditions" : {}
    }
]

```

In this example, only one index scan will be performed. The conditions correspond to the predicates `u.address.state = "CA"` from the query. Specifically, the starting index entry must have the value **CA** on the `address.state` field. All subsequent entries must have **CA** as the value of their `address.state` field, and the scan will stop as soon as an entry with a different state value is encountered. Although the query contains a range predicate on `income`, this predicate does not appear as a range condition of the index scan. This is because there is no equality condition on the `address.city` field that appears before the `income` field in the index definition, and as a result, the `income` predicate cannot be used to determine the boundaries of the scan. Instead, the `income` predicate can be used as an index filtering predicate that is applied on every index entry produced by the index scan.

- The index filtering predicate evaluates the filter criteria on the **income** field. Using the greater than operator the filter condition is evaluated.

```

"index filtering predicate" :
{
    "iterator kind" : "GREATER_THAN",
    "left operand" :
    {
        "iterator kind" : "FIELD_STEP",
        "field name" : "income",
    }
}

```

```

        "input iterator" :
        {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u_idx"
        }
    },
    "right operand" :
    {
        "iterator kind" : "CONST",
        "value" : 5000
    }
}

```

- The index row variable is **\$\$u_idx** which is the name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the **\$\$u_idx** variable will be bound to that entry.
- When the FROM iterator is a TABLE iterator, the FROM variable is the same as either the **index row variable** or the **row variable** of the TABLE iterator, depending on whether the used index is covering or not. In this example, the FROM variable is same as the row variable as the index is not covering.
- This row variable (**\$\$u**) will be referenced by iterators implementing the other clauses of the SELECT expression.
- In the SELECT expression, three fields (**id**, **income** and **age**) are fetched. These correspond to three field names and field expressions in the SELECT expression clause.

```

{
    "field name" : "id",
    "field expression" :
    {
        "iterator kind" : "FIELD_STEP",
        "field name" : "#id",
        "input iterator" :
        {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u"
        }
    }
}

```

For every field to be fetched by the SELECT expression, the field expression is computed by a FIELD_STEP iterator. The FIELD_STEP iterator extracts and returns the value of a field (id as shown above) from the records produced by its input iterator. The same is done for every field to be fetched in the SELECT expression.

Example 4: Sort the data using a Covering index

Example 10-4 Fetch the id and income of users whose state is CA and whose city of residence is Santaclara and whose income is between 1000 and 10000. Sort the result by the income of the users

```

SELECT id, income FROM Users u WHERE u.address.state = "CA" AND
u.address.city= "Santaclara" AND 1000 < income and income < 10000
ORDER BY income

```

Query execution plan:

```

{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_SHARDS",
  "order by fields at positions" : [ 1 ],
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "users",
      "row variable" : "$$u",
      "index used" : "idx_state_city_income",
      "covering index" : true,
      "index row variable" : "$$u_idx",
      "index scans" : [
        {
          "equality conditions" :
          { "address.state": "CA", "address.city": "Santaclara" },
          "range conditions" : { "income" : { "start value" : 1000, "start
inclusive" : false, "end value" : 10000, "end inclusive" : false } }
        }
      ]
    },
    "FROM variable" : "$$u_idx",
    "SELECT expressions" : [
      {
        "field name" : "id",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "#id",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u_idx"
          }
        }
      },
      {
        "field name" : "income",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "income",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u_idx"
          }
        }
      }
    ]
  }
}

```

```
}
}
```

Explanation of the query execution plan :

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator. The only property of the RECEIVE iterator in this example, is the distribution kind whose value is **ALL_SHARDS**.
- The results need to be sorted by income. The income is a part of the **idx_state_city_income** index. So in this example, you don't need a separate SORT operator. The sorting is done by the RECEIVE operator using its property **order by fields at positions**, which is an array. The value of this array depends on the position of the field that is sorted in the SELECT expression.

```
"order by fields at positions" : [ 1 ]
```

In this example, the order by is done using the **income** field, which is the second field in the SELECT expression. That is why you see "1" in the order by fields at position property of the iterator.

Note

If the order of fields in the SELECT expression is different, then the value above changes. For example, if the query is *SELECT income, id FROM Users u WHERE u.address.state = "CA" AND u.address.city= "Santaclara" AND 1000 < income and income < 10000 ORDER BY income*, the order by fields would be **order by fields at positions : [0]** as the **income** field is the first field in the SELECT expression.

- The index **idx_state_city_income** is used here and in this example, it is a covering index as the query can be evaluated using only the entries of the index.
- The index scan property contains the start and stop conditions that define the index scans to be performed.

```
"index scans" : [
  {
    "equality conditions" :
    {"address.state":"CA","address.city":"Santaclara"},
    "range conditions" : { "income" : { "start value" : 1000, "start
inclusive" : false,
                                     "end value" : 10000, "end
inclusive" : false } }
  }
]
```

In this query, only one index scan will be performed. The equality conditions correspond to the predicates `u.address.state = "CA"` and `u.address.city = "Santaclara"` from the query. The index scan will start at the first entry must have the value **CA** on the `address.state` field and its `address.city` field is equal to Santaclara. Specifically, the starting index entry must have the value **CA** on the `address.state` field. All subsequent entries must have **CA** as the value of their `address.state` field, and the scan will stop as soon as an entry with a different state value is encountered. There is a range condition to be applied here on the `income` field.

- The index row variable is **\$\$u_idx** which is the name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the **\$\$u_idx** variable will be bound to that entry.
- When the FROM iterator is a TABLE iterator, the FROM variable is the same as either the **index row variable** or the **row variable** of the TABLE iterator, depending on whether the used index is covering or not. In this example, the FROM variable is same as the index row variable (**\$\$u_idx**) as the index is covering because all the fields in the SELECT expression can be evaluated only using the index entries.
- This index row variable (**\$\$u_idx**) will be referenced by iterators implementing the other clauses of the SELECT expression.
- In the SELECT expression, two fields (**id** and **income**) are fetched. These correspond to two field names and field expressions in the SELECT expression clause.

```
{
  "field name" : "id",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "#id",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u_idx"
    }
  }
}
```

For every field to be fetched by the SELECT expression, the field expression is computed by a FIELD_STEP iterator. The FIELD_STEP iterator extracts and returns the value of a field (id as shown above) from the records produced by its input iterator. The same is done for every field to be fetched in the SELECT expression.

Example 5: Sort the data using a field not part of the index

Example 10-5 Fetch the id, income, and age of users belonging to the state CA and whose city of residence is Santaclara and have income between 1000 and 10000. Sort the results by age

```
SELECT id, income, age FROM Users u WHERE u.address.state = "CA"
AND u.address.city ="Santaclara" AND 1000 < income AND
income < 10000 ORDER BY age
```

Query execution plan:

```
{
  "iterator kind" : "SORT",
  "order by fields at positions" : [ 2 ],
  "input iterator" :
  {
    "iterator kind" : "RECEIVE",
    "distribution kind" : "ALL_SHARDS",
    "input iterator" :
```

```

{
  "iterator kind" : "SELECT",
  "FROM" :
  {
    "iterator kind" : "TABLE",
    "target table" : "users",
    "row variable" : "$$u",
    "index used" : "idx_state_city_income",
    "covering index" : false,
    "index scans" : [
      {
        "equality conditions" :
        {"address.state":"CA","address.city":"Santaclara"},
        "range conditions" : { "income" : { "start value" : 1000, "start
inclusive" : false, "end value" : 10000, "end inclusive" : false } }
      }
    ]
  },
  "FROM variable" : "$$u",
  "SELECT expressions" : [
    {
      "field name" : "id",
      "field expression" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "id",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$u"
        }
      }
    },
    {
      "field name" : "income",
      "field expression" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "income",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$u"
        }
      }
    },
    {
      "field name" : "age",
      "field expression" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "age",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$u"
        }
      }
    }
  ]
}

```

```

    }
  }
}

```

Explanation of the query execution plan :

- In this example, you perform a sort in addition to filtering the data. The results need to be sorted by **age**. The age is not part of the `idx_state_city_income` index. So in this example, you need a separate SORT operator.
- The sorting is done by a **SORT** iterator, which is the parent of the **RECEIVE** iterator. The **order by fields at positions** property specifies the field used for sorting. The value of this array depends on the position of the field that is sorted in the SELECT expression. In this example, age is the third field in the SELECT expression. So **order by fields at positions** has a value of 2.

```
"order by fields at positions" : [ 2 ]
```

- The index `idx_state_city_income` is used here and in this example, it is not a covering index as the query has the age field that is not part of the entries of the index.
- The index scan property contains the start and stop conditions that define the index scans to be performed.

```

"index scans" : [
  {
    "equality conditions" :
    {"address.state":"CA","address.city":"Santaclara"},
    "range conditions" : { "income" : { "start value" : 1000, "start
inclusive" : false,
                                     "end value" : 10000, "end
inclusive" : false } }
  }
]

```

In this example, only one index scan will be performed. The conditions correspond to the predicates `u.address.state = "CA"` and `u.address.city = "Santaclara"` from the query. Specifically, the starting index entry must have the value **CA** on the `address.state` field. All subsequent entries must have **CA** as the value of their `address.state` field, and the scan will stop as soon as an entry with a different state value is encountered. There is a range condition to be applied here on the income field.

- When the FROM iterator is a TABLE iterator, the FROM variable is the same as either the **index row variable** or the **row variable** of the TABLE iterator, depending on whether the used index is covering or not. In this example, the FROM variable is same as the row variable as the index is not covering.
- This row variable (`$$u`) will be referenced by iterators implementing the other clauses of the SELECT expression.

- In the SELECT expression, three fields (**id**, **income** and **age**) are fetched. These correspond to three field names and field expressions in the SELECT expression clause.

```
{
  "field name" : "id",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "#id",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$$u"
    }
  }
}
```

For every field to be fetched by the SELECT expression, the field expression is computed by a FIELD_STEP iterator. The FIELD_STEP iterator extracts and returns the value of a field (id as shown above) from the records produced by its input iterator. The same is done for every field to be fetched in the SELECT expression.

Example 6: Group the data using a Covering index

Example 10-6 Fetch the state, city, and sum of income of all users grouped by the state and city

```
SELECT u.address.state, u.address.city, sum(income)
AS income FROM Users u GROUP BY u.address.state, u.address.city
```

Query execution plan:

```
{
  "iterator kind" : "SELECT",
  "FROM" :
  {
    "iterator kind" : "RECEIVE",
    "distribution kind" : "ALL_SHARDS",
    "order by fields at positions" : [ 0, 1 ],
    "input iterator" :
    {
      "iterator kind" : "SELECT",
      "FROM" :
      {
        "iterator kind" : "TABLE",
        "target table" : "Users",
        "row variable" : "$$u",
        "index used" : "idx_state_city_income",
        "covering index" : true,
        "index row variable" : "$$u_idx",
        "index scans" : [
          {
            "equality conditions" : {},
            "range conditions" : {}
          }
        ]
      }
    }
  }
}
```

```

    }
  ]
},
"FROM variable" : "$$u_idx",
"GROUP BY" : "Grouping by the first 2 expressions in the SELECT
list",
"SELECT expressions" : [
  {
    "field name" : "state",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "address.state",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u_idx"
      }
    }
  },
  {
    "field name" : "city",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "address.city",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$u_idx"
      }
    }
  },
  {
    "field name" : "income",
    "field expression" :
    {
      "iterator kind" : "FUNC_SUM",
      "input iterator" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "income",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$u_idx"
        }
      }
    }
  }
]
},
"FROM variable" : "$from-1",
"GROUP BY" : "Grouping by the first 2 expressions in the SELECT list",
"SELECT expressions" : [

```

```

{
  "field name" : "state",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "state",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$from-1"
    }
  }
},
{
  "field name" : "city",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "city",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$from-1"
    }
  }
},
{
  "field name" : "income",
  "field expression" :
  {
    "iterator kind" : "FUNC_SUM",
    "input iterator" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "income",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$from-1"
      }
    }
  }
}
]
}

```

Explanation of the query execution plan :

- In this example, you group the users based on **state** and **city** and then you determine the **sum** of income of the users grouped.
- The group-by is index-based, that is the group by fields are also part of the index used. This is indicated by the lack of any GROUP iterators. Instead, the grouping is done by the **SELECT** iterators.
- There are two SELECT iterators, the inner one has a **GROUP BY** property that specifies which of the SELECT-clause expressions are also grouping expressions. Here the group

by fields are the first 2 expressions in the SELECT list
(**u.address.state,u.address.city**).

"GROUP BY" : "Grouping by the first 2 expressions in the SELECT list"

- The index **idx_state_city_income** is used here and in this example, it is a covering index as the query can be evaluated using only the entries of the index.
- The index row variable is **\$\$u_idx**, which is the name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the **\$\$u_idx** variable will be bound to that entry.
- This index row variable (**\$\$u_idx**) will be referenced by iterators implementing the other clauses of the SELECT expression.
- In the SELECT expression, three fields (**state, city** and **sum(income)**) are fetched. These correspond to three field names and field expressions in the SELECT expression clause.
- The results returned by the inner SELECT iterators from the various RNs are partial groups, because rows with the same state and city may exist at multiple RNs. So, regrouping and re-aggregation has to be performed at the driver. This is done by the outer SELECT iterator (above the RECEIVE iterator).
- The result is also sorted by **state** and **city**. The **order by fields at positions** property specifies the field used for sorting. The value of this array depends on the position of the field that is sorted in the SELECT expression. In this example, **state** is the first field and **city** is the second field in the SELECT expression. So **order by fields at positions** has a value of 0,1.

"order by fields at positions" : [0, 1]

- In the outer SELECT expression, three fields are fetched: **state, city** and **sum(income)**. The FROM variable **\$from-1** will be referenced by iterators implementing the other clauses of the outer SELECT expression. This corresponds to three field names and field expressions in the outer SELECT expression clause. In this example, two of the field expressions fetch only the fields and there is one field expression which evaluates a function (**sum(income)**).
- For the two fields to be directly fetched by the SELECT expression, the field expression is computed by a FIELD_STEP iterator. The FIELD_STEP iterator extracts and returns the value of a field (state as shown below) from the records produced by its input iterator. The same is done for the city field.

```
{
  "field name" : "state",
  "field expression" :
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "state",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$from-1"
    }
  }
}
```

The third field in the SELECT expression is a function to determine the sum of income. The **FUNC_SUM** iterator is used for this. It iterates over the value of the income field and determines the sum of all incomes from the result of its input iterator.

```
{
  "field name" : "income",
  "field expression" :
  {
    "iterator kind" : "FUNC_SUM",
    "input iterator" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "income",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$from-1"
      }
    }
  }
}
```

Example 7: Group data with fields not part of the index

Example 10-7 Fetch the age and sum of income of all users whose state is CA, grouping the data by age

```
SELECT age, sum(income) FROM Users u
WHERE u.address.state = "CA" GROUP BY age
```

Query execution plan:

```
{
  "iterator kind" : "GROUP",
  "input variable" : "$gb-2",
  "input iterator" :
  {
    "iterator kind" : "RECEIVE",
    "distribution kind" : "ALL_SHARDS",
    "input iterator" :
    {
      "iterator kind" : "GROUP",
      "input variable" : "$gb-1",
      "input iterator" :
      {
        "iterator kind" : "SELECT",
        "FROM" :
        {
          "iterator kind" : "TABLE",
          "target table" : "users",
          "row variable" : "$$u",
          "index used" : "idx_state_city_income",
          "covering index" : false,
          "index scans" : [
```

```

        {
          "equality conditions" : {"address.state":"CA"},
          "range conditions" : {}
        }
      ]
    },
    "FROM variable" : "$$u",
    "SELECT expressions" : [
      {
        "field name" : "age",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "age",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u"
          }
        }
      },
      {
        "field name" : "Column_2",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "income",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$u"
          }
        }
      }
    ]
  },
  "grouping expressions" : [
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "age",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$gb-1"
      }
    }
  ],
  "aggregate functions" : [
    {
      "iterator kind" : "FUNC_SUM",
      "input iterator" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "Column_2",
        "input iterator" :
        {

```

```

        "iterator kind" : "VAR_REF",
        "variable" : "$gb-1"
      }
    }
  ]
},
"grouping expressions" : [
  {
    "iterator kind" : "FIELD_STEP",
    "field name" : "age",
    "input iterator" :
    {
      "iterator kind" : "VAR_REF",
      "variable" : "$gb-2"
    }
  }
],
"aggregate functions" : [
  {
    "iterator kind" : "FUNC_SUM",
    "input iterator" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "Column_2",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$gb-2"
      }
    }
  }
]
}

```

Explanation of the query execution plan :

- In this example , you group all users in the state **CA** based on their age and determine the sum of the income of users belonging to each age group.
- As the GROUP BY field (**age** in this example) is not part of any index, you need a separate GROUP operator to do the grouping. This is indicated by the existence of the GROUP iterators in the execution plan. There are two **GROUP** iterators: one that operates at the driver (above the **RECEIVE** iterator) and another that operates at the RNs (below the **RECEIVE** iterator).
- The lower GROUP iterator has a SELECT iterator as input. The SELECT returns the age and income of each user in the state **CA**. The GROUP iterator will operate until the batch limit is reached. If the batch limit is defined as the max number N of results produced, the GROUP iterator will stop when up to N age groups have been created. If the batch limit is defined as the max number of bytes read, it will stop when this max is reached. The GROUP operator has an input variable. For the inner GROUP operator the input variable is **\$gb-1** and for the outer GROUP operator it is **\$gb-2**.

```

"iterator kind" : "GROUP","input variable" : "$gb-1",

```

- The index **idx_state_city_income** is used here and in this example, it is not a covering index as the query has the **age** field, which is not part of the entries of the index.
- The index scan property contains the start and stop conditions that define the index scans to be performed. In this example, only one index scan will be performed. The conditions correspond to the predicate `u.address.state = "CA"`. Specifically, the starting index entry must have the value **CA** on the `address.state` field. All subsequent entries must have **CA** as the value of their `address.state` field, and the scan will stop as soon as an entry with a different state value is encountered.
- When the FROM iterator is a TABLE iterator, the FROM variable is the same as either the **index row variable** or the **row variable** of the TABLE iterator, depending on whether the used index is covering or not. In this example, the FROM variable is same as the row variable as the index is not covering.
- This row variable (**\$\$u**) will be referenced by iterators implementing the other clauses of the inner SELECT expression.
- The GROUP iterator creates an internal variable (**\$gb-1**) that iterates over the records produced by the SELECT expression.
- The result set produced by the lower GROUP iterator is partial: it may not contain all the age groups and for the age groups that it does contain, the income may be a partial sum (because all rows for a given age may not have been retrieved when query execution stops). The upper GROUP iterator receives the partial results from each RN and performs the final grouping and aggregation. It operates the same way as the lower GROUP iterators and will keep operating until there are no more partial results from the RNs. At that point, the full and final result set is cached at the upper GROUP iterator and is returned to the application.
- The upper GROUP iterator creates an internal variable (**\$gb-2**) that iterates over the records produced by the outer SELECT expression. The \$gb-2 variable has the age and sum of income of all users of the **CA** state, grouped by **age**.
- In the SELECT expression, two fields are fetched: `age` and `sum(income)`. These correspond to two field names and field expressions in the SELECT expression clause.
- For the age field, the field expression is computed by a FIELD_STEP iterator. The FIELD_STEP iterator extracts and returns the value of the age field from the records produced by its input iterator.

```
{
  "iterator kind" : "FIELD_STEP",
  "field name" : "age",
  "input iterator" :
  {
    "iterator kind" : "VAR_REF",
    "variable" : "$gb-2"
  }
}
```

- The second field to be fetched is an aggregate function **sum** of income values. The **FUNC_SUM** iterator is used for this. It iterates over the value of the income field and determines the sum of all incomes from the result of its input iterator.

```
"aggregate functions" : [
  {
    "iterator kind" : "FUNC_SUM",
    "input iterator" :
```



```
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "Column_2",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$gb-2"
      }
    }
  ]
}
```

GeoJson Data Management

This chapter describes GeoJson data and how to search and index GeoJson data in Oracle NoSQL Database. Support for GeoJson data is available only in the Enterprise Edition of Oracle NoSQL Database.

This chapter contains the following topics:

- [About GeoJson Data](#)
- [Lines and Coordinate System](#)
- [Restrictions on GeoJson Data](#)
- [Searching for GeoJson Data](#)
- [Indexing GeoJson Data](#)

About GeoJson Data

The GeoJson specification (Internet Engineering Task Force) defines the structure and content of json objects that are supposed to represent geographical shapes on earth (called geometries). Oracle NoSQL Database implements a number of functions that do indeed interpret such json objects as geometries and allow for the search for rows containing geometries that satisfy certain conditions. Search is made efficient via the use of special indexes.

According to the GeoJson specification, for a json object to be a geometry object it must have two fields called "type" and "coordinates", where the value of the "type" field specifies the kind of geometry and the value of "coordinates" must be an array whose elements define the geometrical shape (the GeometryCollection kind is an exception to this rule, as we will see below). The value of the "type" field must be one of the following 7 strings, corresponding to 7 different kinds of geometry objects: "Point", "LineSegment", "Polygon", "MultiPoint", "MultiLineString", "MultiPolygon", and "GeometryCollection". The value of "coordinates" depends on the kind of geometry, but in all cases it is composed of a number of positions. A position specifies a position on the surface of the earth as an array of 2 double numbers, where the first number is the longitude and the second number is the latitude of the position (GeoJson allows the position's altitude as a 3rd coordinate, but Oracle NoSQL Database does not support altitudes). Longitude and latitude are specified as degrees and must range between -180 to +180 and -90 to +90, respectively.

The 7 kinds of geometry objects are defined as follows: (with an example given in each case)

Point

For type "Point", the "coordinates" field is a single position.

```
{ "type" : "point", "coordinates" : [ 23.549, 35.2908 ] }
```

LineString

A LineString is one or more connected lines; the end-point of one line is the start-point of the next line. The "coordinates" member is an array of two or more positions: the 1st position is

the start point of the 1st line and each subsequent position is the end point of the current line and the start of the next line. Lines may cross each other.

```
{
  "type" : "LineString",
  "coordinates" : [ [121.9447, 37.2975],
    [121.9500, 37.3171],

    [121.9892, 37.3182],
    [122.1554, 37.3882],
    [122.2899, 37.4589],
    [122.4273, 37.6032],
    [122.4304, 37.6267],
    [122.3975, 37.6144]
  ]
}
```

Polygon

A polygon defines a surface area by specifying its outer perimeter and the perimeters of any potential holes inside the area. More precisely, a polygon consists of one or more linear rings, where (a) a linear ring is a closed LineString with four or more positions, (b) the first and last positions are equivalent, and they must contain identical values, (c) a linear ring is the boundary of a surface or the boundary of a hole in a surface, and (d) a linear ring must follow the right-hand rule with respect to the area it bounds, i.e., for exterior rings their positions must be ordered counterclockwise, and for holes their position must be ordered clockwise. Then, the "coordinates" field of a polygon must be an array of linear ring coordinate arrays, where the first must be the exterior ring, and any others must be interior rings. The exterior ring bounds the surface, and the interior rings (if present) bound holes within the surface. The example below shows a polygon with no holes.

```
{
  "type" : "polygon",
  "coordinates" : [ [
    [23.48, 35.16],
    [24.30, 35.16],
    [24.30, 35.50],
    [24.16, 35.61],
    [23.74, 35.70],
    [23.56, 35.60],
    [23.48, 35.16]
  ]
]
```

MultiPoint

For type "MultiPoint", the "coordinates" field is an array of two or more positions.

```
{
  "type" : "MultiPoint",
  "coordinates" : [ [-121.9447, 37.2975],
    [-121.9500, 37.3171],
    [-122.3975, 37.6144]
  ]
}
```

```
]
}
```

MultiLineString

For type "MultiLineString", the "coordinates" member is an array of LineString coordinate arrays.

```
{
  "type": "MultiLineString",
  "coordinates": [
    [ [100.0, 0.0], [01.0, 1.0] ],
    [ [102.0, 2.0], [103.0, 3.0] ]
  ]
}
```

MultiPolygon

For type "MultiPolygon", the "coordinates" member is an array of Polygon coordinate arrays.

```
{
  "type": "MultiPolygon",
  "coordinates": [
    [
      [
        [102.0, 2.0],
        [103.0, 2.0],
        [103.0, 3.0],
        [102.0, 3.0],
        [102.0, 2.0]
      ]
    ],
    [
      [
        [100.0, 0.0],
        [101.0, 0.0],
        [101.0, 1.0],
        [100.0, 1.0],
        [100.0, 0.0]
      ]
    ]
  ]
}
```

GeometryCollection

Instead of a "coordinates" field, a GeometryCollection has a "geometries" field. The value of "geometries" is an array. Each element of this array is a GeoJSON object whose kind is one of the 6 kinds defined earlier. So, in general, a GeometryCollection is a heterogeneous composition of geometries.

```
{
  "type": "GeometryCollection",
  "geometries": [
    {
```

```

"type": "Point",
"coordinates": [100.0, 0.0]
},
{
"type": "LineString",
"coordinates": [ [101.0, 0.0], [102.0, 1.0] ]
}
]
}

```

The GeoJson specification defines 2 additional kinds of entities, called **Feature** and **FeatureCollection**, which allow for combining geometries with other, non-geometrical properties. The specification uses defined above) or a **Feature** or a **FeatureCollection**. **Feature** and **FeatureCollection** are defined as follows:

Feature

A **Feature** object has a "type" member with the value "Feature". A **Feature** object has a "geometry" member, whose value either a geometry object of the 7 kinds defined above or the JSON null value. A **Feature** object has a "properties" member, whose value is any JSON object or the JSON null value.

FeatureCollection

A **FeatureCollection** object has a "type" member with the value "FeatureCollection". A **FeatureCollection** object has a "features" member, whose value is a JSON array. Each element of the array is a **Feature** object as defined above. It is possible for this array to be empty.

Lines and Coordinate System

As shown in the previous section, all kinds of geometries are specified in terms of a set of positions. However, for line strings and polygons, the actual geometrical shape is formed by the lines connecting their positions. The GeoJson specification defines a line between two points as the straight line that connects the points in the (flat) cartesian coordinate system whose horizontal and vertical axes are the longitude and latitude, respectively. More precisely, the coordinates of every point on a line that does not cross the antimeridian between a point $P1 = (lon1, lat1)$ and $P2 = (lon2, lat2)$ can be calculated as:

$$P = (lon, lat) = (lon1 + (lon2 - lon1) * t, lat1 + (lat2 - lat1) * t)$$

with t being a real number greater than or equal to 0 and smaller than or equal to 1.


However, Oracle NoSQL Database uses a geodetic coordinate system (WGS 84) and as a result deviates from the GeoJson specification by using **geodetic lines**: A geodetic line between 2 points is the shortest line that can be drawn between the 2 points on the ellipsoidal surface of the earth. For a simplified, but more illustrative definition, assume for a moment that the earth surface is a sphere. Then, the geodetic line between two points on earth is the minor arc between the two points on the **great circle** corresponding to the points, i.e., the circle that is formed by the intersection of the sphere and the plane defined by the center of the earth and the two points.

The following figure shows the difference between the geodetic and straight lines between Los Angeles and London.


Figure 11-1 Geodetic vs Straight Line

(source: <https://developers.arcgis.com>)

The following figure shows the difference between the two coordinate systems for a square defined by points P1, P2, P3, and P4. The square is assumed to be in the northern hemisphere. The 2 vertical lines of the square are the same in both systems; points on each of these lines have the same longitude. This is not true for the "horizontal" lines. In the GeoJson system all points on the P1-P2 line (the blue line) have the same latitude (so this line is part of an earth parallel). But the geodetic line between P1 and P2 forms a curve (the red line) to the north of the GeoJson line. The difference between the two lines (the curvature) gets more pronounced closer to the poles and as the distance between P1 and P2 increases.

Figure 11-2 Geodetic vs GeoJson Box

When searching for points or other geometries inside the [P1, P2, P3, P4] polygon (using one of the functions described in the next section), Oracle NoSQL Database uses the geodetic view of the search polygon. So, for example, points that are between the blue and the red P1-P2 lines will not appear in the result. What if you really want to search inside the blue box? Such a search can be approximated by adding points between P1-P2 and P4-P3 in the definition of the search polygon. This is illustrated in the following figure, where we have added points P5 and P6. We can see that with the [P1, P6, P2, P3, P5, P4] polygon, the area difference between the geodetic and GeoJson boxes is smaller than with the [P1, P2, P3, P4] polygon.

Figure 11-3 Approximating a Search within a GeoJson Box

Restrictions on GeoJson Data

The following 2 restrictions apply to the kind of GeoJson data supported by Oracle NoSQL Database:

Anti-meridian crossing

Geometries that cross the anti-meridian line cannot be indexed and cannot appear as arguments to the geo search functions described in the following section.

Too big geometries

A geometry is considered "too big" if its Minimum Bounding Box (MBR) has a side whose end points are more than 120 degrees or latitude or longitude apart. Such geometries cannot be indexed and cannot appear as arguments to the geo search functions described in the following section.

Searching for GeoJson Data

Oracle NoSQL Database provides 4 functions to search for GeoJson data that have a certain relationship with a search geometry.

boolean geo_intersect(any*, any*)

Raises an error if it can be detected at compile time that an operand will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns false if any operand returns 0 or more than 1 items.
- Returns NULL if any operand returns NULL.
- Returns false if any operand returns an item that is not a valid geometry object.

Finally, if both operands return a single geometry object, it returns true if the 2 geometries have any points in common; otherwise false.

boolean geo_inside(any*, any*)

Raises an error if it can be detected at compile time that an operand will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns false if any operand returns 0 or more than 1 items.
- Returns NULL if any operand returns NULL.
- Returns false if any operand returns an item that is not a valid geometry object (however, if it can be detected at compile time that an operand will not return a valid geometry, an error is raised).
- Returns false if the second operand returns a geometry object that is not a polygon.

Finally, if both operands return a single geometry object and the second geometry is a polygon, it returns true if the first geometry is completely contained inside the second polygon, i.e., all its points belong to the interior of the polygon; otherwise false. The interior of a polygon is all the points in the polygon area except the points on the linear rings that define the polygon's boundary.

boolean geo_within_distance(any*, any*, double)

Raises an error if it can be detected at compile time that any of the first two operands will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns false if any of the first two operands returns 0 or more than 1 items.
- Returns NULL if any of the first two operands returns NULL.
- Returns false if any of the first two operands returns an item that is not a valid geometry object.

Finally, if both of the first two operands return a single geometry object, it returns true if the first geometry is within a distance of N meters from the second geometry, where N is the number returned by the third operand; otherwise false. The distance between 2 geometries is defined as the minimum among the distances of any pair of points where the first point belongs to the first geometry and the second point to the second geometry. If N is a negative number, it is set to 0.

boolean geo_near(any*, any*, double)

geo_near is converted internally to geo_within_distance plus an (implicit) order-by the distance between the two geometries. However, if the query has an (explicit) order-by already, no ordering by distance is performed. The geo_near function can appear in the WHERE clause only, where it must be a top-level predicate, i.e, not nested under an OR or NOT operator.

In addition to the above search functions, the following two functions are also provided:

double geo_distance(any*, any*)

Raises an error if it can be detected at compile time that an operand will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns -1 if any operand returns zero or more than 1 items.
- Returns NULL if any operand returns NULL.
- Returns -1 if any of the operands is not a geometry.

Otherwise it returns the geodetic distance between the 2 input geometries. The returned distance is the minimum among the distances of any pair of points where the first point belongs to the first geometry and the second point to the second geometry. Between two such points, their distance is the length of the geodetic line that connects the points.

boolean geo_is_geometry(any*)

- Returns false if the operand returns zero or more than 1 items.
- Returns NULL if the operand returns NULL.
- Returns true if the input is a single valid geometry object. Otherwise, false.

Notice that the above geo functions operate on geometry objects, but not on Features or FeatureCollections. Nevertheless, Features and FeatureCollections can still be queried effectively by passing their contained geometry objects to the geo function. An example of this is shown in the following section.

Example 11-1 Searching for GeoJson Data

Consider a table whose rows store points of interest. The table has an id column as its primary key and a poi column of type json.

```
CREATE TABLE PointsOfInterest (
    id INTEGER, poi JSON,
    PRIMARY KEY(id))

INSERT INTO PointsOfInterest VALUES (
    1,
    {
        "kind" : "city hall",
        "address" : {
            "state" : "CA",
            "city" : "Campbell",
            "street" : "70 North 1st street"
        },
        "location" : {
            "type" : "point",
            "coordinates" : [121.94,37.29]
        }
    }
)
INSERT INTO PointsOfInterest VALUES (
    2,
    {
        "kind" : "nature park",
        "name" : "castle rock state park",
        "address" : {
            "state" : "CA",
            "city" : "Los Gatos",
            "street" : "15000 Skyline Blvd"
        },
        "location" : {
            "type" : "polygon",
            "coordinates" : [
                [
                    [122.1301, 37.2330],
                    [122.1136, 37.2256],
                    [122.0920, 37.2291],
                    [122.1020, 37.2347],
                    [122.1217, 37.2380],
                    [122.1301, 37.2330]
                ]
            ]
        }
    }
)
```

The following query looks for nature parks in northern California. The query uses `geo_intersect`, instead of `geo_inside`, to include parks that straddle the border with neighbor states.

```
SELECT t.poi AS park
FROM PointsOfInterest t
WHERE t.poi.kind = "nature park"
AND
geo_intersect(
    t.poi.location,
    {
        "type" : "polygon",
        "coordinates" : [
            [
                [121.94, 36.28],
                [117.52, 37.38],
                [119.99, 39.00],
                [120.00, 41.97],
                [124.21, 41.97],
                [124.39, 40.42],
                [121.94, 36.28]
            ]
        ]
    }
)
```

The following query looks for gas stations within a mile of a given route. The returned gas stations are ordered by ascending distance from the route.

```
SELECT
t.poi AS gas_station,
geo_distance(
    t.poi.location,
    {
        "type" : "LineString",
        "coordinates" : [
            [121.9447, 37.2975],
            [121.9500, 37.3171],
            [121.9892, 37.3182],
            [122.1554, 37.3882],
            [122.2899, 37.4589],
            [122.4273, 37.6032],
            [122.4304, 37.6267],
            [122.3975, 37.6144]
        ]
    }
) AS distance
FROM PointsOfInterest t
WHERE t.poi.kind = "gas station"
AND
geo_near(
    t.poi.location,
    {
        "type" : "LineString",
        "coordinates" : [
```

```

        [121.9447, 37.2975],
        [121.9500, 37.3171],
        [121.9892, 37.3182],
        [122.1554, 37.3882],
        [122.2899, 37.4589],
        [122.4273, 37.6032],
        [122.4304, 37.6267],
        [122.3975, 37.6144]
    ]
},
1609
)

```

Example 11-2 Searching for GeoJson data

This example shows how FeatureCollections can be queried in Oracle NoSQL Database. Consider a "companies" table that stores info about companies, including the locations where each company has offices and some properties for each office location.

```

CREATE TABLE companies (
    id INTEGER, info JSON, PRIMARY KEY(id))

INSERT INTO companies VALUES (
    1,
    {
        "id" : 1,
        "info" : {
            "name" : "acme",
            "CEO" : "some random person",
            "locations" : {
                "type" : "FeatureCollection",
                "features" : [
                    {
                        "type" : "Feature",
                        "geometry" : {
                            "type" : "point",
                            "coordinates" : [ 23.549, 35.2908 ]
                        },
                        "properties" : {
                            "kind" : "development",
                            "city" : "palo alto"
                        }
                    },
                    {
                        "type" : "Feature",
                        "geometry" : {
                            "type" : "point",
                            "coordinates" : [ 23.9, 35.17 ]
                        },
                        "properties" : {
                            "kind" : "sales",
                            "city" : "san jose"
                        }
                    }
                ]
            }
        }
    }
)

```

```

    ]
  }
}
)

```

The following query looks for companies that have sales offices within a search region and returns, for each such company, an array containing the geo-locations of the sales offices within the same search region.

```

SELECT id,
c.info.locations.features [
  geo_intersect(
    $element.geometry,
    {
      "type" : "polygon",
      "coordinates" : [
        [
          [23.48, 35.16],
          [24.30, 35.16],
          [24.30, 35.70],
          [23.48, 35.70],
          [23.48, 35.16]
        ]
      ]
    }
  )
AND
  $element.properties.kind = "sales"
].geometry AS loc
FROM companies c
WHERE EXISTS c.info.locations.features [
  geo_intersect(
    $element.geometry,
    {
      "type" : "polygon",
      "coordinates" : [
        [
          [23.48, 35.16],
          [24.30, 35.16],
          [24.30, 35.70],
          [23.48, 35.70],
          [23.48, 35.16]
        ]
      ]
    }
  )
AND
  $element.properties.kind = "sales"
]

```

For efficient execution of this query, the following index can be created:

```
CREATE INDEX idx_kind_loc ON companies (
  info.locations.features[].properties.kind AS STRING,
  info.locations.features[].geometry AS POINT)
```


Indexing GeoJson Data

Indexing GeoJson data is similar to indexing other json data. In the GeoJson case, the GEOMETRY or POINT keyword must be used after an index path that leads to geometry objects. POINT should be used only if all rows in the table are expected to have single point geometries at the indexed field (GEOMETRY can also be used in this case, but POINT is recommended for better performance). As in the case of other json data, an error will be raised if for some row the value of the index path is not a valid GeoJson point or geometry, unless that value is NULL, json null, or EMPTY.

An index that includes a path to geometry objects is called a geometry index. A geometry index can index other fields as well, but some restrictions apply: (a) a geometry index cannot index more than one GeoJson field, (b) the GeoJson field cannot be inside an array, unless it is a POINT field, and (c) a geometry index cannot be a multi-key index, unless the GeoJson field is a POINT field and the array or map being indexed is the one that contains the POINT field.

Indexing of geometries is based on geohashing. Geohashing is an encoding of a longitude/latitude pair to a string. It works by recursively partitioning the 2-D longitude/latitude coordinate system into a hierarchy of rectangulars called cells. The initial (level-0) cell is the whole world, i.e., all points with a longitude between -180 and +180 and latitude between -90 and +90. The first (level-0) split creates the 32 level-1 cells shown in the following figure. Each cell is assigned a "name", which is a character out of this 32-char-long string G = "0123456789bcdefghjkmnpqrstuvwxy". This name is called the geohash of the cell.

Figure 11-4 32 Level-1 Geohash Cells



The next (level-1) split splits each level-1 cell into 32 level-2 cells. The following figure shows all the level-2 cells inside the "g" cell.

Figure 11-5 Level-1 and Level-2 Geohash Cells



As shown, the geohash of each level-2 is 2 chars long, where the 1st char is the geohash of the parent cell, and the 2nd char is again drawn from the same char set. This process continues down to some given level L, called the geohash length. During an even-numbered split, each cell is split into 8 vertical slices and 4 horizontal slices. During an odd-numbered split, each cell is split into 4 vertical slices and 8 horizontal slices. In both cases, for each of the 32 sub-cells, its geohash is formed by using the parent-cell geohash as a prefix and appending a char out of G. The extra char for each sub-cell is chosen the same way as shown in both the earlier figures for even and odd splits respectively.

Oracle NoSQL Database uses a geohash length of 10. Cells at level 10 have an area of about 1 square meter. When indexing a point, the level-10 cell that contains the point is computed and the geohash of that cell is placed in the index entry. So, for points, a single index entry is generated for each point, and a geometry index on a POINT field behaves like a simple (non-

multikey) index, unless the POINT field itself is inside an array or map that is being indexed. Notice that all points inside the same level-10 cell will have the same geohash.

With the geohashing algorithm described above, points that are close to each other will usually (but not always) be close together in the geometry index as well, i.e., have long common prefixes. So, searching for points using one of the functions described in the previous section translates to one or more range scans in the geometry index. These range scans may return false positives, so the search function itself must still be applied on the rows returned by index scans to eliminate the false positives.

When indexing a LineString or Polygon, the geometry's minimum bounding box (MBR) is computed first, and then a set of cells is found that completely cover the MBR. The level of the covering cells depends on the size and shape and position of the MBR (usually it will be less than 10). Then, an index entry is created for each of the covering cells containing the geohash of that cell. So, a geometry index on a LineString or Polygon is always a multi-key index since multiple index entries will be created for a single row. For MultiPoints, MultiLineStrings, MultiPolygons, and GeometryCollections each of the constituent geometries is indexed separately and the index for such geometries it also a multi-key index.

Example 11-3 Index the coordinates field in a GeoJSON data table

```
CREATE INDEX IF NOT EXISTS CoordIdx ON PointsOfInterest(poi.coordinates AS
    GEOMETRY)
```

Explanation: PointsOfInterest is a [GeoJSON data table](#), which stores geographical points of interest as JSON data in the poi field. The poi field includes the type field, which specifies the kind of geometry and a nested array coordinates field, which includes the geographical coordinates of the points of interest. For more details, see [About GeoJson Data](#).

In the query above, you create an index on the coordinates field. You use the GEOMETRY keyword next to the index path to provide the flexibility to include both single-point and multipoint geometrical types in the coordinates field.

You can verify the index creation as follows:

```
SHOW INDEXES ON PointsOfInterest
```

Output:

```
indexes
  CoordIdx
```

12

Built-in Functions

This chapter discusses about the Built-in functions supported in Oracle NoSQL Database.

Most of the examples demonstrated in this section use one of the two common schemas, that is, the BaggageInfo table schema and stream_acct table schema.

The BaggageInfo table contains passenger details and bag information that is made available as a part of the airline application. The stream_acct table includes information about various shows that the customers watch and can be used on a TV streaming application.

If you want to use the common schema and follow along with the examples, download the scripts `baggageschema_loaddata.sql` and `acctstream_loaddata.sql`.

Start your KVSTORE or KVLite and open the SQL shell.

```
java -jar lib/kvstore.jar kvlite -secure-config disable
```

```
java -jar lib/sql.jar -helper-hosts localhost:5000 -store kvstore
```

Using the `load` command, run the required script.

Example:

```
load -file baggageschema_loaddata.sql
```

```
load -file acctstream_loaddata.sql
```

This creates the tables used in the examples and loads the data into the tables.

For more details on the common schema, see *Getting started with SQL for Oracle NoSQL Database* in the *Developers Guide*.

This chapter contains the following topics:

- [Functions on Complex Values](#)
- [Functions on Sequences](#)
- [Functions on Timestamps](#)
- [Functions on Rows](#)
- [Function to generate a UUID string](#)
- [Functions on GeoJson Data](#)
- [Functions on Strings](#)
- [Function to Convert String to JSON](#)
- [Functions of Mathematical Operations](#)

Functions on Complex Values

size function

Returns the number of fields/entries of a complex item (array, map, record). The function accepts an empty sequence as argument, in which case it will return the empty sequence. The function will return NULL if its input is NULL. The result type is Integer.

Syntax:

```
integer? size(any?)
```

Semantics:

- **any:** The any?element in the above syntax is a complex parameter. The size function accepts any complex parameter such as an array, map, or record. Although the parameter type appears as any?, the function raises an error if the given item is not complex.
- **return type:** Integer

Example 12-1 Determine the number of bags owned by a passenger

Every passenger usually carries one or more pieces of luggage during their travel. In the airline application, you can determine the number of bags owned by a passenger using the size function.

```
SELECT
fullName AS NAME,
size(bagInfo) AS BagCount
FROM BaggageInfo
```

Explanation:

In this example, the information on all the luggage owned by each a passenger is available in the bagInfo array of each passenger record. The size function returns the size of the bagInfo array, which indicates the number of bags. If a passenger owns two pieces of luggage, the bagInfo array has two items and the size function returns the value as 2.

Output:

```
{"NAME":"Adelaide Willard","BagCount":1}
{"NAME":"Raymond Griffin","BagCount":1}
{"NAME":"Henry Jenkins","BagCount":1}
{"NAME":"Lucinda Beckman","BagCount":1}
{"NAME":"Michelle Payne","BagCount":1}
{"NAME":"Joanne Diaz","BagCount":1}
{"NAME":"Mary Watson","BagCount":1}
{"NAME":"Gerard Greene","BagCount":1}
{"NAME":"Fallon Clements","BagCount":1}
{"NAME":"Kendal Biddle","BagCount":1}
{"NAME":"Elane Lemons","BagCount":1}
{"NAME":"Adam Phillips","BagCount":1}
{"NAME":"Lorenzo Phil","BagCount":2}
{"NAME":"Omar Harvey","BagCount":1}
{"NAME":"Lisbeth Wampler","BagCount":1}
```

```
{"NAME":"Dierdre Amador","BagCount":1}  
{"NAME":"Teena Colley","BagCount":1}  
{"NAME":"Rosalia Triplett","BagCount":1}  
{"NAME":"Zulema Martindale","BagCount":1}  
{"NAME":"Doris Martin","BagCount":1}  
{"NAME":"Zina Christenson","BagCount":1}
```

In the query above, since the `bagInfo` array contains one JSON document for each bag checked in by a passenger, the `Bagcount` in this query displays the number of bags that each passenger has checked into their flight.

Functions on Sequences

seq_concat function

`seq_concat` is a variadic function: it can have any number of arguments. It simply evaluates its arguments (if any) in the order they are listed in the argument list, and concatenates the sequences returned by these arguments. This function accepts any data type as the input argument. If the input is a scalar, then the input item is treated as a sequence of size 1.

Syntax:

```
any* seq_concat(any*, ...)
```

Semantics:

- **any:** The *any?* element in the above syntax can be of any data type. The `seq_concat` function accepts any number of input arguments. The arguments have to be in a comma-separated format.
- **return type:** any type

Example 12-2 From the TV streaming application, fetch the shows and genres viewed by users

In the TV streaming application, you can offer suggestions to the users regarding upcoming shows. This is usually based on the shows they have already watched or the genres they prefer.

```
SELECT acct_id,  
  
concat(stream_acct.acct_data[].firstName,  
      ' ',stream_acct.acct_data[].lastName) AS Fullname,  
seq_concat(stream_acct.acct_data[].contentStreamed[].showName) AS Uwatched,  
seq_concat(stream_acct.acct_data[].contentStreamed[].genres) AS Uprefer  
FROM stream_acct  
ORDER BY acct_id
```

Explanation:

In this example, the `seq_concat` function lists all the watched shows and their genres as a comma-separated list.

Output:

```
{
  "acct_id": 1,
  "Fullname": "Adam Phillips",
  "Uwatched": [
    "At the Ranch",
    "Bienvenu"
  ],
  "Uprefer": [
    "action",
    "crime",
    "spanish",
    "comedy",
    "french"
  ]
}

{
  "acct_id": 2,
  "Fullname": "Adelaide Willard",
  "Uwatched": "Bienvenu",
  "Uprefer": [
    "comedy",
    "french"
  ]
}
```

The concat function is also used in the query to concatenate the two strings, firstName and lastName and display it as a single object Fullname in the output. For more details, see [concat Function](#).

Example 12-3 Fetch the details of the shows watched by users

```
SELECT
```

```
    concat(stream_acct.acct_data[].firstName,
  ' ',stream_acct.acct_data[].lastName) AS Fullname,
    CASE
      WHEN exists stream_acct.acct_data[].contentStreamed[].showId
      THEN seq_concat(stream_acct.acct_data[].contentStreamed[0].showName,
stream_acct.acct_data[0].contentStreamed[0].seriesInfo[0].episodes[0],
stream_acct.acct_data[0].contentStreamed[0].seriesInfo[0].episodes[1],stream_a
cct.acct_data[].contentStreamed[1].showName,
stream_acct.acct_data[0].contentStreamed[0].seriesInfo[1].episodes[0],
stream_acct.acct_data[0].contentStreamed[0].seriesInfo[1].episodes[1])
      ELSE "Start streaming your favorite shows here"
    END AS Showdetails
FROM stream_acct WHERE acct_id=1
```

Explanation:

If a user has watched any show, the details are stored in the contentSreamed JSON field. The above query retrieves the date of streaming, details of the episode, the duration, and the time elapsed for each show. The seq_concat function is used to concatenate and display all the details in the output.

Output:

```
{
  "Fullname": "Adam Phillips",
  "Showdetails": [
    "At the Ranch",
    {
      "date": "2022-04-18",
      "episodeID": 20,
      "lengthMin": 85,
      "minWatched": 85
    },
    {
      "date": "2022-04-18",
      "episodeID": 30,
      "lengthMin": 60,
      "minWatched": 60
    },
    "Bienvenu",
    {
      "date": "2022-04-25",
      "episodeID": 40,
      "lengthMin": 50,
      "minWatched": 50
    },
    {
      "date": "2022-04-27",
      "episodeID": 50,
      "lengthMin": 45,
      "minWatched": 30
    }
  ]
}
```

The concat function is also used in the query to concatenate the two strings, firstName and lastName and display it as a single object Fullname in the output. For more details, see [concat Function](#).

In addition to the above there are also the following aggregate functions on sequences. They are described in the [Sequence Aggregate Functions](#) section.

- long seq_count(any*)
- number seq_sum(any*)

- number seq_avg(any*)
- any_atomic seq_min(any*)
- any_atomic seq_max(any*)

seq_distinct function

Returns the distinct values of its input sequence.

Syntax:

any* seq_distinct(any*)

Semantics:

- **any:** The seq_distinct function accepts any parameter type as the input argument.
- **return type:** any type

Example 12-4 Determine the number of users having phone numbers in different area codes

Consider an application that maintains the information of the users. Create the schema for users table as follows:

```
CREATE TABLE users ( id INTEGER,
                      firstName STRING,
                      lastName STRING,
                      otherNames ARRAY(RECORD(first STRING, last STRING)),
                      age INTEGER,
                      income INTEGER,
                      address JSON,
                      connections ARRAY(INTEGER),
                      expenses MAP(INTEGER),
PRIMARY KEY (id) )
```

Insert the following user record into the table.

```
INSERT INTO users VALUES (
    10,
    "John",
    "Smith",
    [ {"first" : "Johny", "last" : "BeGood"} ],
    22,
    45000,
    {
        "street" : "Pacific Ave",
        "number" : 101,
        "city" : "Santa Cruz",
        "state" : "CA",
        "zip" : 95008,
        "phones" : [
            { "area" : 408, "number" : 4538955, "kind" : "work" },
            { "area" : 831, "number" : 7533341, "kind" : "home" },
            { "area" : 831, "number" : 7533382, "kind" : "mobile" }
        ]
    },
    {}
),
```

```
[ 30, 55, 43 ],
DEFAULT
)
```

Explanation:

A user may have multiple phone numbers in the same area code. To determine the number of users having phone numbers in different area codes, a particular user should be counted only once (even if the user has more than one phone number with the same area code).

This is possible using the `seq_distinct` function, which eliminates duplicate values from its input sequence.

```
SELECT $area, count(*) AS cnt
FROM Users u, seq_distinct(u.address.phones.area) AS $area
GROUP BY $area
```

Output:

```
{"area":408,"cnt":1}
{"area":831,"cnt":1}
```

The user John Smith has three phone numbers, two of which have the same area code. However, only distinct area codes are fetched in the query.

Functions on Timestamps

The timestamp functions perform various operations on the supplied timestamps.

You can add a duration to a timestamp, find the difference between two timestamps, and round timestamp to a specified unit. You can cast a timestamp to/from string with customized patterns. Some of the functions support the extraction of the date part of a timestamp. You can also use these functions to display the current time.

A few of these timestamp functions take a second argument, which specifies the units or the format patterns.

The input timestamp can be a [TIMESTAMP](#) value/expression. If required, the timestamp functions perform an implicit [CAST](#) on the supplied timestamp.

You can invoke the timestamp functions from the `SELECT/WHERE` clauses, and also from other SQL clauses where function calls are allowed in the syntax. For example, you can supply the functions as arguments to [Aggregate Functions](#).

The following timestamp functions are supported:

Table 12-1 Timestamp functions

Function	Description
timestamp_add	Adds a duration to a timestamp value.
timestamp_diff	Returns the number of milliseconds between two timestamp values.
get_duration	Converts the given number of milliseconds to a duration string.

Table 12-1 (Cont.) Timestamp functions

Function	Description
timestamp_ceil	Rounds-up the timestamp value to the specified unit.
timestamp_floor/timestamp_trunc	Rounds-down the timestamp value to the specified unit.
timestamp_round	Rounds the timestamp value to the specified unit.
timestamp_bucket	Rounds the timestamp value to the beginning of the specified interval, starting from a specified origin value.
format_timestamp	Converts a timestamp into a string according to the specified pattern and the timezone.
parse_to_timestamp	Converts a string in the specified pattern into a timestamp value.
to_last_day_of_month	Returns the last day of the month from a given timestamp.
Timestamp extract functions	Extracts the corresponding date part of a given timestamp. The following functions are supported: • year • month • day • hour • minute • second • millisecond • microsecond • nanosecond Returns the week number within the year. The following functions are supported: • week • isoweek Returns the corresponding index from a given timestamp. The following functions are supported: • quarter • day_of_week • day_of_month • day_of_year
current_time_millis	Returns the current time as the number of milliseconds.
current_time	Returns the current time as a timestamp value.

The examples in the following sections are based on an airline baggage tracking application and a TV streaming application. To follow along with the examples, create the tables and load data as described in [Tables used in the examples](#) section.

timestamp_add function

Adds a duration to a timestamp value and returns the new timestamp. The duration can be positive or negative. The result type is `TIMESTAMP(9)`.

Syntax:

```
TIMESTAMP(9) timestamp_add(TIMESTAMP timestamp, STRING duration)
```

Semantics:

- **timestamp:** A [TIMESTAMP](#) value or a value that can be cast to `TIMESTAMP`.
- **duration:** A [STRING](#) with format `[-](<n> <UNIT>)+`, where 'n' is a number and the `<UNIT>` can be `YEAR`, `MONTH`, `DAY`, `HOURL`, `MINUTE`, `SECOND`, `MILLISECOND`, `NANOSECOND` or the plural form of these keywords (e.g. `YEARS`).

Note

The `UNIT` keyword is case-insensitive.

- **Return Value:** `TIMESTAMP(9)`

Example 12-5 In the airline application, a buffer of five minutes delay is considered "on-time" . Print the estimated arrival time on the first leg with a buffer of five minutes for the passenger with ticket number 1762399766476

```
SELECT timestamp_add(bag.bagInfo.flightLegs[0].estimatedArrival, "5 minutes")
AS ARRIVAL_TIME FROM BaggageInfo bag WHERE ticketNo=1762399766476
```

Explanation: In the airline application, a customer can have any number of flight legs depending on the source and destination. In the query above, you are fetching the estimated arrival in the "first leg" of the travel. So the first record of the `flightLeg` array is fetched and the `estimatedArrival` time is fetched from the array and a buffer of "5 minutes" is added to that and displayed.

Output:

```
{"ARRIVAL_TIME":"2019-02-03T06:05:00.000000000Z"}
```

Note

The column `estimatedArrival` is a `STRING`. If the column has `STRING` values in ISO-8601 format, then it will be automatically converted by the SQL runtime into `TIMESTAMP` datatype.

Example 12-6 Print the estimated arrival time in every leg with a buffer of five minutes for the passenger with ticket number 1762399766476

```
SELECT $s.ticketno, $value as estimate, timestamp_add($value, '5 minute') AS
add5min
FROM baggageinfo $s,$s.bagInfo.flightLegs.estimatedArrival as $value
WHERE ticketNo=1762399766476
```

Explanation: You want to display the `estimatedArrival` time in every leg. The number of legs can be different for every customer. So variable reference is used in the query above and the `baggageInfo` array and the `flightLegs` array are unnested to execute the query.

Output:

```
{"ticketno":1762399766476,"estimate":"2019-02-03T06:00:00Z","add5min":"2019-02-03T06:05:00.000000000Z"}
```

```
{"ticketno":1762399766476,"estimate":"2019-02-03T08:22:00Z","add5min":"2019-02-03T08:27:00.000000000Z"}
```

Example 12-7 Find the number of bags arrived in the last week

```
SELECT count(*) AS COUNT_LASTWEEK FROM baggageInfo bag WHERE  
EXISTS bag.bagInfo[$element.bagArrivalDate > current_time()  
AND $element.bagArrivalDate < timestamp_add(current_time(), "-7 days")]
```

Explanation: You get a count of the number of bags processed by the airline application in the last week. A customer can have more than one bag(that is bagInfo array can have more than one record). The bagArrivalDate should have a value between today and the last 7 days. For every record in the bagInfo array, you determine if the bag arrival time is between the time now and one week ago. The function `current_time` gives you the time now. An EXISTS condition is used as a filter for determining if the bag has an arrival date in the last one week. The count function determines the total number of bags in this time period.

Example 12-8 Find the number of bags arriving in the next 6 hours

```
SELECT count(*) AS COUNT_NEXT6HOURS FROM baggageInfo bag WHERE  
exists bag.bagInfo[$element.bagArrivalDate > current_time()  
AND $element.bagArrivalDate < timestamp_add(current_time(), "6 hours")]
```

Explanation: You get a count of the number of bags that will be processed by the airline application in the next 6 hours. A customer can have more than one bag(that is bagInfo array can have more than one record). The bagArrivalDate should be between the time now and the next 6 hours. For every record in the bagInfo array, you determine if the bag arrival time is between the time now and six hours later. The function `current_time` gives you the time now. An EXISTS condition is used as a filter for determining if the bag has an arrival date in the next six hours. The count function determines the total number of bags in this time period.

timestamp_diff function

The `timestamp_diff` function returns the number of milliseconds between two timestamp values. The result type is LONG.

Syntax:

```
LONG timestamp_diff(TIMESTAMP timestamp1, TIMESTAMP timestamp2
```

Semantics:

- **timestamp1:** A TIMESTAMP value or a value that can be cast to TIMESTAMP
- **timestamp2:** A TIMESTAMP value or a value that can be cast to TIMESTAMP
- **Return Value:** LONG

Example 12-9 What is the duration between the time the baggage was boarded at one leg and reached the next leg for the passenger with ticket number 1762355527825?

To determine the duration in milliseconds, use the `timestamp_diff` function.

```
SELECT $bagInfo.bagArrivalDate, $flightLeg.flightDate,  
timestamp_diff($bagInfo.bagArrivalDate, $flightLeg.flightDate) AS diff
```



```
FROM baggageinfo $s,
$s.bagInfo[] AS $bagInfo, $bagInfo.flightLegs[] AS $flightLeg
WHERE ticketNo=176235527825
```

Explanation: In an airline application every customer can have different number of hops/legs between their source and destination. In this query, you determine the time taken between every flight leg. This is determined by the difference between bagArrivalDate and flightDate for every flight leg.

Output:

```
{"bagArrivalDate":"2019-03-22T10:17:00Z","flightDate":"2019-03-22T07:00:00Z",
"diff":11820000}
{"bagArrivalDate":"2019-03-22T10:17:00Z","flightDate":"2019-03-22T07:23:00Z",
"diff":10440000}
{"bagArrivalDate":"2019-03-22T10:17:00Z","flightDate":"2019-03-22T08:23:00Z",
"diff":6840000}
```

Example 12-10 How long does it take from the time of check-in to the time the bag is scanned at the point of boarding for the passenger with ticket number 176234463813?

To determine the duration in milliseconds, use the timestamp_diff function.

```
SELECT $flightLeg.flightNo,
$flightLeg.actions[contains($element.actionCode, "Checkin")].actionTime AS
checkinTime,
$flightLeg.actions[contains($element.actionCode, "BagTag Scan")].actionTime
AS bagScanTime,
timestamp_diff(
    $flightLeg.actions[contains($element.actionCode, "Checkin")].actionTime,
    $flightLeg.actions[contains($element.actionCode, "BagTag Scan")].actionTime
) AS diff
FROM baggageinfo $s,
$s.bagInfo[].flightLegs[] AS $flightLeg
WHERE ticketNo=176234463813 AND
starts_with($s.bagInfo[].routing, $flightLeg.fltRouteSrc)
```

Explanation: In the baggage data, every flightLeg has an actions array. There are three different actions in the actions array. The action code for the first element in the array is Checkin/Offload. For the first leg, the action code is Checkin and for the other legs, the action code is Offload at the hop. The action code for the second element of the array is BagTag Scan. In the query above, you determine the difference in action time between the bag tag scan and check-in time. You use the contains function to filter the action time only if the action code is Checkin or BagScan. Since only the first flight leg has details of check-in and bag scan, you additionally filter the data using starts_with function to fetch only the source code fltRouteSrc.

Output:

```
{"flightNo":"BM572","checkinTime":"2019-03-02T03:28:00Z","bagScanTime":"2019-03-02T04:52:00Z","diff":-5040000}
```

Example 12-11 How long does it take for the bags of a customer with ticket no 1762320369957 to reach the first transit point?

To determine the duration in milliseconds, use the `timestamp_diff` function.

```
SELECT $bagInfo.flightLegs[0].flightDate,  
$bagInfo.flightLegs[0].estimatedArrival,  
timestamp_diff($bagInfo.flightLegs[0].estimatedArrival,  
$bagInfo.flightLegs[0].flightDate) AS diff  
FROM baggageinfo $s, $s.bagInfo[] AS $bagInfo  
WHERE ticketNo=1762320369957
```

Explanation: In an airline application every customer can have different number of hops/legs between their source and destination. In the example above, you determine the time taken for the bag to reach the first transit point. In the baggage data, the `flightLeg` is an array. The first record in the array refers to the first transit point details. The `flightDate` in the first record is the time when the bag leaves the source and the `estimatedArrival` in the first flight leg record indicates the time it reaches the first transit point. The difference between the two gives the time taken for the bag to reach the first transit point.

Output:

```
{"flightDate":"2019-03-12T03:00:00Z","estimatedArrival":"2019-03-12T16:00:00Z"  
,"diff":46800000}  
{"flightDate":"2019-03-12T03:00:00Z","estimatedArrival":"2019-03-12T16:40:00Z"  
,"diff":49200000}
```

get_duration function

The `get_duration` function converts the given number of milliseconds to a duration string. The result type is `STRING`.

Syntax:

```
STRING get_duration(LONG duration_millis)
```

Semantics:

- **duration_millis:** the duration in milliseconds.
- **Return Value:** `STRING`. The returned duration string format is `[-](<n> <UNIT>)+`, where the `<UNIT>` can be `DAY`, `HOUR`, `MINUTE`, `SECOND` and `MILLISECOND`, e.g. "1 day 2 hours" or "-10 minutes 0 second 500 milliseconds".

Example 12-12 What is the duration in days, hours, or minutes between the time the baggage was boarded at one leg and reached the next leg for the passenger with ticket number 1762355527825?

```
SELECT $s.ticketno, $bagInfo.bagArrivalDate, $flightLeg.flightDate,  
  
get_duration(timestamp_diff($bagInfo.bagArrivalDate, $flightLeg.flightDate))  
AS diff FROM  
    baggageinfo $s, $s.bagInfo[] AS $bagInfo, $bagInfo.flightLegs[]  
AS $flightLeg WHERE  
    ticketNo=1762355527825;
```

Explanation: In an airline application every customer can have different number of hops/legs between their source and destination. In this query, you determine the time taken between every flight leg. This is determined by the difference between bagArrivalDate and flightDate for every flight leg.

The bagArrivalDate and flightDate is in an unsupported format for the get_duration function, so wrap it in the [timestamp_diff function](#) to make it valid.

Output:

```
{ "bagArrivalDate": "2019-03-22T10:17:00Z", "flightDate": "2019-03-22T07:00:00Z", "diff": "3 hours 17 minutes" }
{ "bagArrivalDate": "2019-03-22T10:17:00Z", "flightDate": "2019-03-22T07:23:00Z", "diff": "2 hours 54 minutes" }
{ "bagArrivalDate": "2019-03-22T10:17:00Z", "flightDate": "2019-03-22T08:23:00Z", "diff": "1 hour 54 minutes" }
```

Example 12-13 How long does it take in days, hours, or minutes from the time of check-in to when the bag is scanned at the point of boarding for the passenger with ticket number 176234463813?

```
SELECT $flightLeg.flightNo, $flightLeg.actions[contains($element.actionCode,
    "Checkin")].actionTime AS
    checkinTime, $flightLeg.actions[contains($element.actionCode,
    "BagTag Scan")].actionTime AS bagScanTime, get_duration(timestamp_diff(
    $flightLeg.actions[contains($element.actionCode, "Checkin")].actionTime,
    $flightLeg.actions[contains($element.actionCode, "BagTag
    Scan")].actionTime )) AS diff FROM
    baggageinfo $s, $s.bagInfo[].flightLegs[] AS $flightLeg WHERE
    ticketNo=176234463813 AND
    starts_with($s.bagInfo[].routing, $flightLeg.fltRouteSrc)
```

Explanation: In the baggage data, every flightLeg has an actions array. There are three different actions in the actions array. The action code for the first element in the array is Checkin/Offload. For the first leg, the action code is Checkin and for the other legs, the action code is Offload at the hop. The action code for the second element of the array is BagTag Scan. In the query above, you determine the difference in action time between the bag tag scan and check-in time. You use the contains function to filter the action time only if the action code is Checkin or BagScan. Since only the first flight leg has details of check-in and bag scan, you additionally filter the data using starts_with function to fetch only the source code fltRouteSrc.

The input values for the get_duration function is in an unsupported format, so pass it in the [timestamp_diff function](#) to make it valid.

Output:

```
{ "flightNo": "BM572", "checkinTime": "2019-03-02T03:28:00Z",
  "bagScanTime": "2019-03-02T04:52:00Z", "diff": "- 1 hour 24 minutes" }
```

Example 12-14 How long does it take in days, hours, or minutes for the baggage of a customer with ticket no 1762320369957 to reach the first transit point?

```
SELECT $bagInfo.flightLegs[1].actions[2].actionTime,
    $bagInfo.flightLegs[0].actions[0].actionTime,
```

```
get_duration(timestamp_diff($bagInfo.flightLegs[1].actions[2].actionTime,
$bagInfo.flightLegs[0].actions[0].actionTime)) AS diff FROM
baggageinfo $s, $s.bagInfo[] as $bagInfo
WHERE ticketNo=1762320369957;
```

Explanation: In an airline application every customer can have different number of hops/legs between their source and destination. In the example above, you determine the time taken for the bag to reach the first transit point. In the baggage data, the `flightLeg` is an array. The first record in the array refers to the first transit point details. The `flightDate` in the first record is the time when the bag leaves the source and the `estimatedArrival` in the first flight leg record indicates the time it reaches the first transit point. The difference between the two gives the time taken for the bag to reach the first transit point.

The input values for the `get_durrantion` function is in an unsupported format, so pass it in the [timestamp_diff function](#) to make it valid.

Output:

```
{"flightDate":"2019-03-12T03:00:00Z","estimatedArrival":"2019-03-12T16:00:00Z",
"diff":"13 hours"}
{"flightDate":"2019-03-12T03:00:00Z","estimatedArrival":"2019-03-12T16:40:00Z",
"diff":"13 hours 40 minutes"}
```

timestamp_ceil function

The `timestamp_ceil` function returns the rounded-up value of the given timestamp to the specified unit.

If the input timestamp value is already rounded up to the specified unit, then the return value is the same as the input timestamp value.

Syntax:

```
TIMESTAMP timestamp_ceil(<timestamp>[, unit])
```

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to `TIMESTAMP` type.
- **unit:** The unit argument is optional and a [STRING](#) data type. If not specified, `DAY` is the default unit. For more details, see [Supported units](#).
- **Return Value:** `TIMESTAMP(0)`
The function returns `NULL` in the following cases:
 - If either the timestamp or unit argument is set to `NULL`.
 - If the input timestamp is not castable to `TIMESTAMP` type.

Example 12-15 For airline passengers with reservation code 'LE6J4Z' to 'ZG8Z5N', print the routing information and bag arrival dates rounded up to closest hour

```
SELECT $b.routing AS ROUTE,
timestamp_ceil($b.bagArrivalDate, 'HOUR') AS BAGTIME
FROM BaggageInfo bag, bag.bagInfo[0] AS $b
WHERE confNo BETWEEN 'LE6J4Z' and 'ZG8Z5N'
```

Explanation: Use the `timestamp_ceil` function with the unit value of `HOURL` to round up the checked bag arrival dates to the beginning of the next hour.

To avoid the duplication of results due to multiple checked baggage by a passenger, you consider only the first element of the `bagInfo` array in this query.

Output:

```
{ "ROUTE": "SFO/IST/ATH/JTR", "BAGTIME": "2019-02-03T00:00:00Z" }
{ "ROUTE": "JFK/IST/VIE", "BAGTIME": "2019-03-05T12:00:00Z" }
{ "ROUTE": "SFO/IST/ATH/JTR", "BAGTIME": "2019-03-12T16:00:00Z" }
{ "ROUTE": "MSQ/FRA/HKG", "BAGTIME": "2019-02-03T09:00:00Z" }
{ "ROUTE": "MXP/CDG/SLC/BZN", "BAGTIME": "2019-03-15T11:00:00Z" }
{ "ROUTE": "MIA/LAX/MEL", "BAGTIME": "2019-02-04T11:00:00Z" }
{ "ROUTE": "SFO/ORD/FRA", "BAGTIME": "2019-03-02T14:00:00Z" }
{ "ROUTE": "SFO/IST/ATH/JTR", "BAGTIME": "2019-03-12T16:00:00Z" }
{ "ROUTE": "SFO/IST/ATH/JTR", "BAGTIME": "2019-03-07T17:00:00Z" }
{ "ROUTE": "MIA/LAX/MEL", "BAGTIME": "2019-02-02T00:00:00Z" }
{ "ROUTE": "MEL/LAX/MIA", "BAGTIME": "2019-03-02T17:00:00Z" }
{ "ROUTE": "MXP/CDG/SLC/BZN", "BAGTIME": "2019-02-21T15:00:00Z" }
{ "ROUTE": "MSQ/FRA/HKG", "BAGTIME": "2019-02-13T12:00:00Z" }
{ "ROUTE": "JFK/MAD", "BAGTIME": "2019-03-07T14:00:00Z" }
```

Example 12-16 From airline baggage tracking data, print the bag arrival date and the bag auction date for a passenger with ticket number 1762344493810, considering 90 days as the luggage retention period

```
SELECT $b.bagArrivalDate AS BagArrival,
timestamp_ceil(timestamp_add($b.bagArrivalDate, '90 Days'), 'day') AS
BagCollection
FROM BaggageInfo bag, bag.bagInfo AS $b
WHERE ticketNo=1762344493810
```

Explanation: This query shows how to nest the timestamp functions. To determine the date an unclaimed bag is retained, add 90 days to the `bagArrivalDate` using the [timestamp_add](#) function. The `timestamp_ceil` function rounds up the value to the beginning of the next day.

Output:

```
{ "BagArrival": "2019-02-01T16:13:00Z", "BagCollection": "2019-05-03T00:00:00Z" }
```

timestamp_floor or timestamp_trunc function

The `timestamp_floor` or `timestamp_trunc` function returns the rounded-down value of the given timestamp to the specified unit. The functions can be used interchangeably in a query.

If the input timestamp value is already rounded down to the specified unit, then the return value is the same as the input timestamp value.

Syntax:

```
TIMESTAMP timestamp_floor(<timestamp>[, unit])
TIMESTAMP timestamp_trunc(<timestamp>[, unit])
```

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to TIMESTAMP type.
- **unit:** The unit argument is optional and a [STRING](#) data type. If not specified, DAY is the default unit. For more details, see [Supported units](#).
- **Return Value:** TIMESTAMP(0)
The function returns NULL in the following cases:
 - If either the timestamp or unit argument is set to NULL.
 - If the input timestamp is not castable to TIMESTAMP type.

Example 12-17 Print the name, flight number, and travel date for all the passengers who boarded at originating airport JFK in the month of March 2019

```
SELECT bag.fullName, $f.flightNo, $f.flightDate
FROM BaggageInfo bag, bag.bagInfo[0].flightLegs[0] AS $f
WHERE $f.fltRouteSrc = "JFK" AND timestamp_floor($f.flightDate, 'MONTH') =
'2019-03-01'
```

Explanation: You use the `timestamp_floor` function with the unit value as MONTH to round down the travel dates to the beginning of the month. You then compare the resulting timestamp value with the string "2019-03-01" to select the desired passengers. This query does not consider the passengers in transit.

This example supplies the date in an ISO-8601 formatted string, which gets implicitly CAST into a TIMESTAMP value.

To avoid the duplication of results due to multiple checked bags by a passenger, you consider only the first element of the `bagInfo` array in this query.

Output:

```
{"fullName":"Kendal
Biddle","flightNo":"BM127","flightDate":"2019-03-04T06:00:00Z"}
{"fullName":"Dierdre
Amador","flightNo":"BM495","flightDate":"2019-03-07T07:00:00Z"}
```

Example 12-18 Fetch the flight number, flight departure time, and security check-in time for a passenger

```
SELECT $b.flightLegs[0].flightNo,
$b.flightLegs[0].flightDate,
timestamp_add(timestamp_trunc($b.flightLegs[0].flightDate, 'HOUR'), '-2
HOURS') AS SECURITYCHECK
FROM BaggageInfo bag, bag.bagInfo[0] AS $b
WHERE ticketNo=1762344493810
```

Explanation: Usually in an airline, the passengers are allowed to proceed through security check approximately two hours before the departure. To calculate the security check-in time, you subtract two hours from the flight departure time. You first approximate the `flightDate` value by using the `timestamp_trunc` function with HOUR as unit to round down the `flightDate` to the beginning of the hour. Supply the resulting timestamp to the [timestamp_add](#) function with a negative value of '-2 HOURS' to subtract the truncated flight departure time by two hours.

To avoid the duplication of results due to multiple checked baggage by a passenger, you consider only the first element of the `bagInfo` array in this query.

Output:

```
{"flightNo":"BM604","flightDate":"2019-02-01T06:00:00Z","SECURITYCHECK":"2019-02-01T04:00:00.000000000Z"}
```

timestamp_round function

The `timestamp_round` function returns a timestamp after rounding the supplied timestamp to the specified unit.

Note

See also, [timestamp_ceil](#) and [timestamp_floor/timestamp_trunc](#) functions.

If the input timestamp value is already rounded to the specified unit, then the return value is the same as the input timestamp value.

Syntax:

```
TIMESTAMP timestamp_round(<timestamp>[, unit])
```

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to `TIMESTAMP` type.
- **unit:** The unit argument is optional and a [STRING](#) data type. If not specified, `DAY` is the default unit. For more details, see [Supported units](#).
- **Return Value:** `TIMESTAMP(0)`
The function returns `NULL` in the following cases:
 - If either the timestamp or unit argument is set to `NULL`.
 - If the input timestamp is not castable to `TIMESTAMP` type.

The units and their behavior are as follows:

Table 12-2 timestamp_round functional behavior

Unit	Actions
YEAR, IYEAR	Rounds up to the next year if the timestamp is on or after July first, otherwise rounds down.
QUARTER	Rounds up to the next quarter if the timestamp is on or after the 16th day of the 2nd month of the quarter, otherwise rounds down.
MONTH	Rounds up to the next month if the timestamp is on or after the 16th day, otherwise rounds down.
WEEK, IWEEK	Rounds up to the next week if the timestamp is on or after the mid-day(12 PM) of the 4th day of the week, otherwise rounds down.

Table 12-2 (Cont.) timestamp_round functional behavior

Unit	Actions
DAY	Rounds up to the next day if the timestamp is on or after the mid-day(12 PM) of the current day, otherwise rounds down.
HOURL	Rounds up to the next hour if the timestamp is on or after half an hour of the current hour, otherwise rounds down.
MINUTE	Rounds up to the next minute if the timestamp is on or after 30 seconds of the current minute, otherwise rounds down.
SECOND	Rounds up to the next second if the timestamp is on or after 500 milliseconds of the current second, otherwise rounds down.

Example 12-19 From the airline baggage tracking data, print all the activities performed on the checked bags in the originating station MEL. Align the actions to one minute interval

```
SELECT $b.actionAt,
       $b.actionCode,
       timestamp_round($b.actionTime, 'MINUTE') as actionTime
FROM baggageInfo bag, bag.bagInfo[0].flightLegs[0].actions[] AS $b
WHERE bag.bagInfo[0].flightLegs[0].fltRouteSrc = "MEL"
```

Explanation: In this query, you use the timestamp_round function with unit as MINUTE to round the actionTime to the nearest minute.

To avoid the duplication of results due to multiple checked baggage by a passenger, you consider only the first element of the bagInfo array in this query.

Output:

```
{"actionAt":"MEL","actionCode":"ONLOAD to
LAX","actionTime":"2019-03-01T12:20:00Z"}
{"actionAt":"MEL","actionCode":"BagTag Scan at
MEL","actionTime":"2019-03-01T11:52:00Z"}
{"actionAt":"MEL","actionCode":"Checkin at
MEL","actionTime":"2019-03-01T11:43:00Z"}
```

timestamp_bucket function

The timestamp_bucket function rounds the given timestamp value to the beginning of the specified interval (bucket). The interval starts at a specified origin on the timeline.

You can use this function for aggregating time series data to a desired time interval, known as periodicity. In certain cases, it is desirable to place all your time series data into equidistant buckets of given periodicity, with each bucket representing the same amount of time.

Syntax:

```
TIMESTAMP timestamp_bucket(<timestamp>[, interval [,origin ]])
```

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to TIMESTAMP type.
- **interval:** The interval argument is optional and a [STRING](#) data type. The interval is specified as <n> unit.
where,
n specifies the value of the interval. The n must be > 0
unit defines the interval component. The function supports WEEK, DAY, HOUR, MINUTE, and SECOND in either singular or plural format.
For example, "5 MINUTE" or "5 MINUTES".

Note

The units are not case-sensitive.

- **origin:** The origin argument represents the starting point of buckets on the timeline. This argument is optional and takes a TIMESTAMP value. The origin can be of any data type that can be cast to TIMESTAMP type. If not specified, Unix epoch 1970-01-01 is the default value.

Note

The function also rounds the input timestamps that are lesser than the origin to the beginning of the specified interval. That is, you can supply an origin with a future timestamp value as compared to the input timestamp value on the timeline.

- **Return Value:** TIMESTAMP(9)
The function returns NULL in the following cases:
 - If any of the arguments are set to NULL.
 - If the input timestamp is not castable to TIMESTAMP type.

Example 12-20 Fetch the statistics of the number of passengers departing from the IST airport every 12 hrs with buckets starting from January 1st, 2019. Consider data only for the month of February 2019

```
SELECT $t AS DATE,
count($t) AS FLIGHTCOUNT
FROM BaggageInfo bag, bag.bagInfo[0].flightLegs[] $f,
timestamp_bucket($f.flightDate, '12 HOURS', '2019-01-01T00') $t
WHERE $f.fltRouteSrc =any "IST" AND timestamp_floor($f.flightDate, 'MONTH') =
'2019-02-01T00:00:00Z'
GROUP BY $t
ORDER BY $t
```

Explanation: To consider passengers traveling in February 2019, use the [timestamp_floor](#) function and round down the flightDate to the beginning of the month. Compare the result with the string "2019-02-01T00:00:00Z". This example supplies the date in an ISO-8601 formatted string, which gets implicitly CAST into a TIMESTAMP value.

To include the transit flights from the IST airport, use the array constructor [] to indicate that the flightLegs is an array and consider each fltRouteSrc array element in the search.

Use the `timestamp_bucket` function on the `flightDate` fields with interval as 12 hours and origin as 1st of January 2019.

Output:

```
{ "DATE": "2019-02-02T12:00:00.000000000Z", "FLIGHTCOUNT": 1 }
{ "DATE": "2019-02-04T00:00:00.000000000Z", "FLIGHTCOUNT": 1 }
{ "DATE": "2019-02-04T12:00:00.000000000Z", "FLIGHTCOUNT": 2 }
{ "DATE": "2019-02-07T12:00:00.000000000Z", "FLIGHTCOUNT": 1 }
{ "DATE": "2019-02-11T12:00:00.000000000Z", "FLIGHTCOUNT": 1 }
{ "DATE": "2019-02-12T00:00:00.000000000Z", "FLIGHTCOUNT": 2 }
{ "DATE": "2019-02-12T12:00:00.000000000Z", "FLIGHTCOUNT": 1 }
```

Example 12-21 From the Streaming media service data, fetch the statistics of number of accounts that expire each week with buckets starting from noon of January 1st, 2023. Consider data from 1st of December 2023

```
SELECT array_collect(s.account_expiry) as ACCOUNT_EXPIRY,
$t AS WEEK,
count($t) AS ACCOUNTS
FROM stream_acct s,
timestamp_bucket(s.account_expiry, '1 week', '2023-01-01T12') $t
WHERE s.account_expiry >= '2023-12-01'
GROUP BY $t
ORDER BY $t;
```

Explanation: Use the `timestamp_bucket` function on the `account_expiry` fields with interval as 1 week and origin as 1st of January 12 PM. The function calculates a periodicity of a week from 1st of January 2023, 12 PM and rounds the `account_expiry` values to the beginning of the corresponding week bucket.

Notice that you use the [array_collect](#) function to display the `account_expiry` fields. You can reference only grouping expressions, aggregate functions, or external variables when the `SELECT` expression includes a grouping function.

Output:

```
{ "ACCOUNT_EXPIRY":
[ "2023-12-18T00:00:00.000000000Z" ], "WEEK": "2023-12-17T12:00:00.000000000Z", "AC
COUNTS": 1 }
{ "ACCOUNT_EXPIRY":
[ "2023-12-31T00:00:00.000000000Z" ], "WEEK": "2023-12-24T12:00:00.000000000Z", "AC
COUNTS": 1 }
{ "ACCOUNT_EXPIRY":
[ "2024-03-18T00:00:00.000000000Z" ], "WEEK": "2024-03-17T12:00:00.000000000Z", "AC
COUNTS": 1 }
```

format_timestamp function

The `format_timestamp` function converts a timestamp into a string according to the specified pattern and the timezone.

Syntax:

```
STRING format_timestamp(<timestamp>,[pattern [, timezone]])
```

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to a TIMESTAMP type.
- **pattern:** The pattern argument is optional and takes [STRING](#) data type as an input. It supports all pattern symbols in Java DateTimeFormatter class, except the timezone symbols 'z', 'zz', 'zzz', and 'v'. For more details on which timezone symbols are supported, see the table below:

Symbol	Meaning	Presentation	Example
V	time-zone ID	zone-id	America/Los_Angeles; Z; -08:30
O	localized zone-offset	Offset-O	GMT+8; GMT+08:00; UTC-08:00
X	zone-offset 'Z' for zero	offset-X	Z; -08; -0830; -08:30; -083015; -08:30:15
x	zone-offset	offset-x	+0000; -08; -0830; -08:30; -083015; -08:30:15
Z	zone-offset	offset-Z	+0000; -0800; -08:00

Note

The default pattern is ISO-8601 format: **yyyy-MM-dd'T'HH:mm:ss[.S..S]**.

- **timezone:** The timezone argument is optional and takes [STRING](#) data type as an input. The timezone argument uses TimeZoneID (an identifier that represents a specific timezone). For example, use well-defined name such as "Asia/Calcutta", or a custom ID such as "GMT-08:00". For more examples, see List of TimeZoneID. Default is UTC.

Note

Except for UTC and GMT, use the well-defined names for the timezones instead of abbreviations (for example, PST, IST).

- **Return Value:** STRING
The function returns NULL in the following cases:
 - If the timestamp, pattern, or timezone argument is set to NULL.
 - If the input timestamp is not castable to TIMESTAMP type.

Example 12-22 For a passenger with a specific ticket number, print the estimated arrival time on the first leg according to the pattern and the timezone entered.

```
SELECT $info.estimatedArrival, format_timestamp($info.estimatedArrival, "MMM
dd, yyyy HH:mm:ss 0", "America/Vancouver") AS FormattedTimestamp
FROM BaggageInfo bag, bag.bagInfo.flightLegs[0] AS $info
WHERE ticketNo= 1762399766476
```

Explanation: In this query, you specify the `estimatedArrival` field, pattern, and full name of the timezone as arguments to the `format_timestamp` function to convert the timestamp string to the specified "MMM dd, yyyy HH:mm:ss" pattern.

Note

The letter 'O' in the pattern argument represents the `ZoneOffset`, which prints the amount of time that differs from Greenwich/UTC in the resulting string.

Output:

```
{"estimatedArrival":"2019-02-03T06:00:00Z","FormattedTimestamp":"Feb 02, 2019  
22:00:00 GMT-8"}
```

Example 12-23 Display the formatted timestamp.

```
SELECT format_timestamp('2024-05-08T09:41:00',"dd MMM, uuuu HH:mm:ss","Asia/  
Calcutta") AS TIMESTAMP1  
FROM BaggageInfo  
WHERE ticketNo=1762399766476
```

Explanation: In this query, you use the `format_timestamp` function to print the given timestamp in the specified pattern. During formatting, the function adjust the timestamp to the "Asia/Calcutta" timezone and displays the resulting timestamp in the "dd MMM, uuuu HH:mm:ss" pattern.

Output:

```
{"TIMESTAMP1":"08 May, 2024 15:11:00"}
```

parse_to_timestamp function

The `parse_to_timestamp` function converts a string into a timestamp, requiring the pattern parameter to match the format of the specified string.

Syntax:

```
TIMESTAMP parse_to_timestamp(<string>[, pattern])
```

Semantics:

- string:** The string argument takes a [STRING](#) data type of a [TIMESTAMP](#) value in the format of the specified pattern.

Note

The string argument with the timestamp uses `TimeZoneID` (an identifier that represents a specific timezone). Except for UTC and GMT, use the well-defined names for the timezones instead of abbreviations (for example, PST, IST).

- **pattern:** The pattern argument is optional and takes [STRING](#) data type as an input. It supports all pattern symbols in Java DateTimeFormatter class, except the timezone symbols 'z', 'zz', 'zzz', and 'v'. For more details on which timezone symbols are supported, see the table below.

Symbol	Meaning	Presentation	Example
V	time-zone ID	zone-id	America/Los_Angeles; Z; -08:30
O	localized zone-offset	Offset-O	GMT+8; GMT+08:00; UTC-08:00
X	zone-offset 'Z' for zero	offset-X	Z; -08; -0830; -08:30; -083015; -08:30:15
x	zone-offset	offset-x	+0000; -08; -0830; -08:30; -083015; -08:30:15
Z	zone-offset	offset-Z	+0000; -0800; -08:00

Note

The default pattern is ISO-8601 format: **yyyy-MM-dd'T'HH:mm:ss[.S..S]**

The table below displays what the pattern argument needs to contain to match a string argument containing timezone information.

If string argument	Then pattern argument must include
ends with a 'Z'	zone symbols like, 'VV', 'X', 'zzzz'.
ends with a full timezone name, for example, America/Los_Angeles, Europe/Paris, and so on. For more example, see List of TimeZoneID.	zone symbols like, 'VV', 'zzzz'.
ends with GMT +/- some value	zone symbols like, 'VV', 'zzzz' or ZoneOffset 'O', 'OOOO'.

- **Return Value:** `TIMESTAMP(9)`
 - The function returns NULL if the string or pattern argument is NULL.

Example 12-24 Parse the streaming account's expiry date string for a user using a specified pattern.

```
SELECT parse_to_timestamp(sa.account_expiry, "yyyy-MM-
dd'T'hh:mm:ss.SSSSSSSSzzzz")AS DAY
FROM stream_acct sa
WHERE acct_id=1
```

Explanation: In the query, you must supply a pattern that matches the format of the string argument, and the `parse_to_timestamp` function will convert the `account_expiry` date into a timestamp value.

Output:

```
{"DAY": "2023-10-18T00:00:00.000000000Z"}
```

Example 12-25 Parse the given string with the specified pattern, which includes a zone symbol, into a timestamp.

```
SELECT parse_to_timestamp('03/14/23 09:04:01 America/Los_Angeles', "MM/dd/yy  
HH:mm:ss VV") AS TIMESTAMP  
FROM BaggageInfo  
WHERE ticketNo=1762390789239
```

Explanation: In this query, use the zone symbol, 'VV', for any TimeZoneID in the string argument. Unless wrapped in a [format_timestamp function](#), the output timestamp will be displayed relative to UTC.

Output:

```
{"TIMESTAMP": "2023-03-14T16:04:01.000000000Z"}
```

Example 12-26 Parse the given string with the specified pattern, which includes a zone offset, into a timestamp.

```
SELECT format_timestamp(parse_to_timestamp('2024/02/12 18:30:54 GMT+02:00',  
"yyyy/dd/MM HH:mm:ss 0000"), "yyyy-MM-dd HH:mm:ss 0000", "GMT+02:00") AS  
TIMESTAMP  
FROM BaggageInfo  
WHERE ticketNo=1762390789239
```

Explanation: In this query, the string argument has a TimeZoneID, GMT+02:00, so the pattern argument must include a zone symbol or a ZoneOffset. When wrapped in the [format_timestamp function](#), the output timestamp will display in the GMT+02:00 timezone.

Output:

```
{"TIMESTAMP": "2024-12-02 18:30:54 GMT+02:00"}
```

to_last_day_of_month function

The `to_last_day_of_month` function returns the last day of the month for a given timestamp.

Syntax:

```
TIMESTAMP to_last_day_of_month(<timestamp>)
```

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to a `TIMESTAMP` type
- **Return Value:** `TIMESTAMP(0)`
The function return NULL in the following cases:
 - If the timestamp argument is NULL.
 - If the input timestamp is not castable to `TIMESTAMP` type.

Example 12-27 For a subscriber, print the last day of the month in which the account subscription expires.

```
SELECT sa.acct_id, to_last_day_of_month(sa.account_expiry) AS lastday
FROM stream_acct sa
WHERE profile_name="DM"
```

Output:

```
{"acct_id":4, "lastday":"2024-03-31T00:00:00Z"}
```

Timestamp extract functions

The timestamp extract functions allow you to fetch the date/week number/timestamp index from the input timestamps. The following sections describe these functions.

- [year/month/day/hour/minute/second/millisecond/microsecond/nanosecond functions](#)
- [week/isoweek functions](#)
- [quarter/week/month/year index functions](#)

Date extract functions

These functions return the corresponding year/month/day/hour/minute/second/millisecond/microsecond/nanosecond from a timestamp. The timestamp must be supplied as an argument to the function.

Syntax:

INTEGER year(<timestamp>)

INTEGER month(<timestamp>)

INTEGER day(<timestamp>)

INTEGER hour(<timestamp>)

INTEGER minute(<timestamp>)

INTEGER second(<timestamp>)

INTEGER millisecond(<timestamp>)

INTEGER microsecond(<timestamp>)

INTEGER nanosecond(<timestamp>)

Semantics:

- **timestamp:** These functions expects a timestamp as the input argument.
- **Return Value:** INTEGER

Table 12-3 Timestamp date extract functions

Function	Return Value
year	Returns the year for the given timestamp. The returned value is in the range -6383 to 9999.
month	Returns the month for the given timestamp, in the range 1 ~ 12.
day	Returns the day of month for the timestamp, in the range 1 ~ 31.
hour	Returns the hour of day for the timestamp, in the range 0 ~ 23.
minute	Returns the minute for the timestamp, in the range 0 ~ 59.
second	Returns the second for the timestamp, in the range 0 ~ 59.
millisecond	Returns the fractional second in millisecond for the timestamp, in the range 0 ~ 999.
microsecond	Returns the fractional second in microsecond for the timestamp, in the range 0 ~ 999999.
nanosecond	Returns the fractional second in nanosecond for the timestamp, in the range 0 ~ 999999999.

Note

If the argument is NULL or empty, the result is also NULL or empty.

Example 12-28 Get consolidated travel details of the passengers from airline baggage tracking data

In an airline application, it is beneficial to the passengers to have a quick summary of their upcoming travel details. You can use miscellaneous time functions to get consolidated travel details of the passengers from the BaggageInfo table.

```
SELECT DISTINCT
  $s.fullName,
  $s.bagInfo[].flightLegs[].flightNo AS flightnumbers,
  $s.bagInfo[].flightLegs[].fltRouteSrc AS From,
  concat ($t1,":", $t2,":", $t3) AS Traveldate
FROM baggageinfo $s, $s.bagInfo[].flightLegs[].flightDate AS $bagInfo,
day(CAST($bagInfo AS Timestamp(0))) $t1,
month(CAST($bagInfo AS Timestamp(0))) $t2,
year(CAST($bagInfo AS Timestamp(0))) $t3
```

Explanation:

You can use the time functions to retrieve the travel date, month, and year. The [concat](#) function is used to concatenate the retrieved travel records to display them in the desired format on the application. You first use [CAST expression](#) to convert the flightDates to a **TIMESTAMP** and then fetch the date, month, and year details from the timestamp.

Output:

```
{"fullName":"Adam Phillips","flightnumbers":["BM604","BM667"],"From":
["MIA","LAX"],"Traveldate":"1:2:2019"}
```

```
{"fullName":"Adelaide Willard","flightnumbers":["BM79","BM907"],"From":
["GRU","ORD"],"Traveldate":"15:2:2019"}
```


The query returns the flight details which can serve as a quick look-up for the passengers.

week/isoweek functions

These functions return the corresponding week/isoweek from a timestamp. The timestamp must be supplied as an argument to the function.

Syntax:

```
INTEGER week(<timestamp>)
```

```
INTEGER isoweek(<timestamp>)
```

Semantics:

- **timestamp:** These functions expects a timestamp as the input argument.
- **Return Value:** INTEGER

Table 12-4 Timestamp week extract functions

Function	Return Value
week	Returns the week number within the year where a week starts on Sunday and the first week has a minimum of 1 day in this year, in the range 1 ~ 54.
isoweek	Returns the week number within the year based on ISO-8601, where a week starts on Monday and the first week has a minimum of 4 days in this year, in range 0 ~ 53.

Note

If the argument is NULL or empty, the result is also NULL or empty.

Example 12-29 Determine the week and ISO week number from a passenger's travel date

```
SELECT
$s.fullName,
$s.contactPhone,
week(CAST($bagInfo.flightLegs[1].flightDate AS Timestamp(0))) AS TravelWeek,
isoweek(CAST($bagInfo.flightLegs[1].flightDate AS Timestamp(0))) AS
ISO_TravelWeek
FROM baggageinfo $s, $s.bagInfo[] AS $bagInfo
```

Explanation: You first use [CAST expression](#) to convert the flightDate to a TIMESTAMP and then fetch the week and isoweek from the timestamp.

Output:

```
{"fullName": "Adelaide
Willard", "contactPhone": "421-272-8082", "TravelWeek": 7, "ISO_TravelWeek": 7}
```

```
{"fullName": "Adam  
Phillips", "contactPhone": "893-324-1064", "TravelWeek": 5, "ISO_TravelWeek": 5}
```

Timestamp index extract functions

These functions return the corresponding quarter/week/month/year index from a timestamp. The timestamp must be supplied as an argument to the function.

Syntax:

INTEGER quarter(<timestamp>)

INTEGER day_of_week(<timestamp>)

INTEGER day_of_month(<timestamp>)

INTEGER day_of_year(<timestamp>)

Semantics:

- **timestamp:** The timestamp argument takes a [TIMESTAMP](#) value or a value that can be cast to TIMESTAMP type. If you supply an unsupported timestamp format, the functions return an error.
- **Return Value:** INTEGER
The functions adhere to the ISO-8601 calendar system.

Table 12-5 Timestamp index extraction functions

Function	Return Value
quarter	1 to 4 1 indicates Quarter 1, and 4 indicates Quarter 4.
day_of_week	1 to 7 1 indicates the first day of the week (Monday) and 7 is the last day of the week (Sunday).
day_of_month	1 to 31 1 indicates the first day of the month and 31 is the last day of the month.
day_of_year	1 to 366 1 indicates the first day of the year and 366 is the last day of the year.

Note

The functions return NULL in the following cases:

- If the timestamp argument is set to NULL.
- If the input timestamp is not castable to TIMESTAMP type.

Example 12-30 Print the quarter in which a subscriber viewed the given episode from a series

```
SELECT
quarter(show.acct_data.contentStreamed[1].seriesInfo[0].episodes[1].date) AS
quarter
FROM stream_acct show
WHERE profile_name="AP"
```

Output:

```
{"quarter":1}
```

Example 12-31 Find the day of the week for given timestamps

```
SELECT day_of_week("2024-06-19") AS DAYVAL1,
day_of_week(parse_to_timestamp('06/19/24', 'MM/dd/yy')) AS DAYVAL2
FROM BaggageInfo
WHERE ticketNo=1762344493810
```

Explanation: The second timestamp in the query is in an unsupported format '06/19/24' by itself, so wrap it in the [parse_to_timestamp](#) function to make it valid.

Output:

```
{
  "DAYVAL1" : 3,
  "DAYVAL2" : 3
}
```

Example 12-32 From the airline baggage tracking data, print the flight numbers, bag arrival date, and the day on which the checked bags arrive at MEL airport

```
SELECT DISTINCT $b.flightLegs[].flightNo,
$b.bagArrivalDate,
day_of_week($b.bagArrivalDate) AS DAY
FROM BaggageInfo bag, bag.bagInfo[0] AS $b
WHERE regex_like($b.routing, ".*MEL")
```

Explanation: You use the `day_of_week` function to determine the day index corresponding to the timestamp value in the `bagArrivalDate` field.

To display the results for the MEL airport, use the [regular expression condition](#) to match the pattern of the destination airport with the string "MEL".

Output:

```
{"flightNo":["BM604","BM667"],"bagArrivalDate":"2019-02-25T20:15:00Z","DAY":1}
{"flightNo":["BM604","BM667"],"bagArrivalDate":"2019-02-04T10:08:00Z","DAY":1}
{"flightNo":["BM604","BM667"],"bagArrivalDate":"2019-02-16T16:13:00Z","DAY":6}
{"flightNo":["BM604","BM667"],"bagArrivalDate":"2019-02-01T16:13:00Z","DAY":5}
```

Example 12-33 From the airline baggage tracking data, print the flight numbers, departure date, and the day of the month for all the flights starting from SFO airport

```
SELECT DISTINCT $f.flightNo AS FLIGHT,  
$f.flightDate AS DEPARTURE,  
day_of_month($f.flightDate) AS DAY  
FROM BaggageInfo bag, bag.bagInfo[0].flightLegs[0] AS $f  
WHERE $f.fltRouteSrc =any "SFO"
```

Output:

```
{"FLIGHT": "BM318", "DEPARTURE": "2019-02-02T12:00:00Z", "DAY": 2}  
{"FLIGHT": "BM318", "DEPARTURE": "2019-03-12T03:00:00Z", "DAY": 12}  
{"FLIGHT": "BM572", "DEPARTURE": "2019-03-02T05:00:00Z", "DAY": 2}  
{"FLIGHT": "BM318", "DEPARTURE": "2019-03-07T04:00:00Z", "DAY": 7}
```

Example 12-34 From the airline baggage tracking data, find the passengers who traveled on the 60th day of the launch of flight "BM114"

```
SELECT fullName, $f.fltRouteSrc AS BOARDING,  
$f.flightNo AS FLIGHT,  
$f.flightDate AS DEPARTURE  
FROM BaggageInfo bag, bag.bagInfo[0].flightLegs[0] AS $f  
WHERE day_of_year($f.flightDate)= 60 AND $f.flightNo =any "BM114"
```

Output:

```
{"fullName": "Omar  
Harvey", "BOARDING": "MEL", "FLIGHT": "BM114", "DEPARTURE": "2019-03-01T12:00:00Z"}
```

current_time_millis function

Returns the current time in UTC, as the number of milliseconds since January 1, 1970 UTC.

Syntax:

```
long current_time_millis()
```

Semantics:

- This function does not expect any input argument.
- **Return Value:** long

Example: See [current_time function](#).

current_time function

Returns the current time in UTC, as a timestamp value with millisecond precision.

Syntax:

```
timestamp(3) current_time()
```

Semantics:

- This function does not expect any input argument.
- **Return Value:** timestamp(3)

Example 2: Miscellaneous timestamp functions**Example 12-35 Determine the time lapse between the last travel date of a passenger and the current date**

In an airline application, a few customers travel very frequently and are entitled to frequent flier miles rewards. You can determine the time lapse between the last travel date of a passenger and the current date to assess if they can be considered for such a reward program.

```
SELECT
$s.fullName,
$s.contactPhone,
get_duration(timestamp_diff(current_time(),
CAST($bagInfo.flightLegs[1].flightDate AS Timestamp(0)))) AS LastTravel
FROM baggageinfo $s, $s.bagInfo[] AS $bagInfo
```

Explanation:

You can use the `current_time` function to get the current time. To determine the timespan between the last travel date and the current date, you can supply the current time to the `get_duration/timestamp_diff` function along with the last travel time. For more details on `timestamp_diff` and `get_duration` functions, see [timestamp_diff function](#).

Output:

```
{"fullName":"Adelaide
Willard","contactPhone":"421-272-8082","LastTravel":"1453 days 6 hours 20
minutes 56 seconds 601 milliseconds"}

{"fullName":"Adam Phillips","contactPhone":"893-324-1064","LastTravel":"1451
days 23 hours 19 minutes 39 seconds 543 milliseconds"}
```

You use the `current_time` function to calculate the current time. Use the `timestamp_diff` function to calculate the time difference between the current time and the last flight date. You first use [CAST expression](#) to convert the flightDates to a `TIMESTAMP` and then fetch the day, month, and year details from the timestamp. Since the `timestamp_diff` function returns the number of milliseconds between two timestamp values, you then use the `get_duration` function to convert the milliseconds to a duration string.

The `get_duration` function converts the milliseconds to days, hours, minutes, seconds, and milliseconds based on the return value. The following conversions are considered for calculation purposes:

```
1000 milliseconds = 1 second
60 seconds = 1 minute
60 minutes = 1 hour
24 hours = 1 day
```

For example: If the `timestamp_diff` function returns the value **129084684821** milliseconds, the `get_duration` function converts it correspondingly to **1494 days 52 minutes 4 seconds 687 milliseconds**.

Supported units

The `unit` is an identifier that specifies the precision to be considered while rounding the input timestamp. The unit takes a [STRING](#) value. You can supply a unit as the second argument to the following timestamp functions: **`timestamp_ceil`**, **`timestamp_floor`**, **`timestamp_round`**, and **`timestamp_trunc`**.

You can specify the following units. If not specified, the unit defaults to `DAY`.

 Note

The units are not case-sensitive.

Table 12-6 Description and Examples of Units

Units	Description
YEAR, IYEAR	- YEAR: The first day of a year. - IYEAR: The Monday of the first calendar week as defined by the ISO 8601 standard. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'YEAR') AS CEIL1, timestamp_ceil('2019-10-21T14:16:00Z', 'YEAR') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'YEAR') AS FLOOR1, timestamp_floor('2019-10-21T14:16:00Z', 'YEAR') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'YEAR') AS ROUND1, timestamp_round('2019-10-21T14:16:00Z', 'YEAR') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'YEAR') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:00Z', 'YEAR') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2020-01-01T00:00:00Z", "CEIL2" : "2020-01-01T00:00:00Z", "FLOOR1" : "2019-01-01T00:00:00Z", "FLOOR2" : "2019-01-01T00:00:00Z", "ROUND1" : "2019-01-01T00:00:00Z", "ROUND2" : "2020-01-01T00:00:00Z", "TRUNC1" : "2019-01-01T00:00:00Z", "TRUNC2" : "2019-01-01T00:00:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
QUARTER	The first day of the quarter as defined by the ISO 8601 standard. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'QUARTER') AS CEIL1, timestamp_ceil('2019-10-21T14:16:00Z', 'QUARTER') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'QUARTER') AS FLOOR1, timestamp_floor('2019-10-21T14:16:00Z', 'QUARTER') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'QUARTER') AS ROUND1, timestamp_round('2019-11-21T14:16:00Z', 'QUARTER') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'QUARTER') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:00Z', 'QUARTER') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-07-01T00:00:00Z", "CEIL2" : "2020-01-01T00:00:00Z", "FLOOR1" : "2019-04-01T00:00:00Z", "FLOOR2" : "2019-10-01T00:00:00Z", "ROUND1" : "2019-04-01T00:00:00Z", "ROUND2" : "2020-01-01T00:00:00Z", "TRUNC1" : "2019-04-01T00:00:00Z", "TRUNC2" : "2019-10-01T00:00:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
MONTH	The first day of the month. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'MONTH') AS CEIL1, timestamp_ceil('2019-10-21T14:16:00Z', 'MONTH') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'MONTH') AS FLOOR1, timestamp_floor('2019-10-21T14:16:00Z', 'MONTH') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'MONTH') AS ROUND1, timestamp_round('2019-11-21T14:16:00Z', 'MONTH') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'MONTH') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:00Z', 'MONTH') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-06-01T00:00:00Z", "CEIL2" : "2019-11-01T00:00:00Z", "FLOOR1" : "2019-05-01T00:00:00Z", "FLOOR2" : "2019-10-01T00:00:00Z", "ROUND1" : "2019-05-01T00:00:00Z", "ROUND2" : "2019-12-01T00:00:00Z", "TRUNC1" : "2019-05-01T00:00:00Z", "TRUNC2" : "2019-10-01T00:00:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
WEEK,IWEEK	- WEEK: Same day of the week as 1st of January. For example: If the first day of the year in the given timestamp is a Tuesday, using the WEEK unit rounds the timestamp up or down to the nearest Tuesday depending on the timestamp function used. - IWEEK: The first day of the calendar week as defined by the ISO 8601 standard, which is Monday. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'WEEK') AS CEIL1, timestamp_ceil('2019-10-21T14:16:00Z', 'WEEK') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'WEEK') AS FLOOR1, timestamp_floor('2019-10-21T14:16:00Z', 'WEEK') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'WEEK') AS ROUND1, timestamp_round('2019-11-15T05:00:00Z', 'WEEK') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'WEEK') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:00Z', 'WEEK') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-05-21T00:00:00Z", "CEIL2" : "2019-10-22T00:00:00Z", "FLOOR1" : "2019-05-14T00:00:00Z", "FLOOR2" : "2019-10-15T00:00:00Z", "ROUND1" : "2019-05-14T00:00:00Z", "ROUND2" : "2019-11-12T00:00:00Z", "TRUNC1" : "2019-05-14T00:00:00Z", "TRUNC2" : "2019-10-15T00:00:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
DAY	Day. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'DAY') AS CEIL1, timestamp_ceil('2019-10-21T14:16:00Z', 'DAY') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'DAY') AS FLOOR1, timestamp_floor('2019-10-21T14:16:00Z', 'DAY') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'DAY') AS ROUND1, timestamp_round('2019-11-15T15:00:00Z', 'DAY') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'DAY') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:00Z', 'DAY') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-05-15T00:00:00Z", "CEIL2" : "2019-10-22T00:00:00Z", "FLOOR1" : "2019-05-14T00:00:00Z", "FLOOR2" : "2019-10-21T00:00:00Z", "ROUND1" : "2019-05-14T00:00:00Z", "ROUND2" : "2019-11-16T00:00:00Z", "TRUNC1" : "2019-05-14T00:00:00Z", "TRUNC2" : "2019-10-21T00:00:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
HOURL	Hour. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'HOURL') AS CEIL1, timestamp_ceil('2019-10-21T14:16:00Z', 'HOURL') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'HOURL') AS FLOOR1, timestamp_floor('2019-10-21T14:16:00Z', 'HOURL') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'HOURL') AS ROUND1, timestamp_round('2019-11-15T15:30:00Z', 'HOURL') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'HOURL') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:00Z', 'HOURL') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-05-14T12:00:00Z", "CEIL2" : "2019-10-21T15:00:00Z", "FLOOR1" : "2019-05-14T11:00:00Z", "FLOOR2" : "2019-10-21T14:00:00Z", "ROUND1" : "2019-05-14T12:00:00Z", "ROUND2" : "2019-11-15T16:00:00Z", "TRUNC1" : "2019-05-14T11:00:00Z", "TRUNC2" : "2019-10-21T14:00:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
MINUTE	Minute. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.000000000Z', 'MINUTE') AS CEIL1, timestamp_ceil('2019-10-21T14:16:24Z', 'MINUTE') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.000000000Z', 'MINUTE') AS FLOOR1, timestamp_floor('2019-10-21T14:16:24Z', 'MINUTE') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.000000000Z', 'MINUTE') AS ROUND1, timestamp_round('2019-11-15T15:30:24Z', 'MINUTE') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.000000000Z', 'MINUTE') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:24Z', 'MINUTE') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-05-14T11:48:00Z", "CEIL2" : "2019-10-21T14:17:00Z", "FLOOR1" : "2019-05-14T11:47:00Z", "FLOOR2" : "2019-10-21T14:16:00Z", "ROUND1" : "2019-05-14T11:48:00Z", "ROUND2" : "2019-11-15T15:30:00Z", "TRUNC1" : "2019-05-14T11:47:00Z", "TRUNC2" : "2019-10-21T14:16:00Z" } </pre>

Table 12-6 (Cont.) Description and Examples of Units

Units	Description
SECOND	Second. Example: <pre> SELECT timestamp_ceil('2019-05-14T11:47:56.4999Z', 'SECOND') AS CEIL1, timestamp_ceil('2019-10-21T14:16:24.99999Z', 'SECOND') AS CEIL2, timestamp_floor('2019-05-14T11:47:56.6999Z', 'SECOND') AS FLOOR1, timestamp_floor('2019-10-21T14:16:24.500Z', 'SECOND') AS FLOOR2, timestamp_round('2019-05-14T11:47:56.4999Z', 'SECOND') AS ROUND1, timestamp_round('2019-11-15T15:30:24.6999Z', 'SECOND') AS ROUND2, timestamp_trunc('2019-05-14T11:47:56.355Z', 'SECOND') AS TRUNC1, timestamp_trunc('2019-10-21T14:16:24.7000Z', 'SECOND') AS TRUNC2 FROM BaggageInfo WHERE ticketNo=1762344493810 </pre> Output: <pre> { "CEIL1" : "2019-05-14T11:47:57Z", "CEIL2" : "2019-10-21T14:16:25Z", "FLOOR1" : "2019-05-14T11:47:56Z", "FLOOR2" : "2019-10-21T14:16:24Z", "ROUND1" : "2019-05-14T11:47:56Z", "ROUND2" : "2019-11-15T15:30:25Z", "TRUNC1" : "2019-05-14T11:47:56Z", "TRUNC2" : "2019-10-21T14:16:24Z" } </pre>

Function to generate a UUID string

The function `random_uuid` returns a randomly generated UUID, as a string of 36 characters. This function can be used to generate values for columns of type UUID in an INSERT or UPDATE SQL statement.

Syntax:

```
string random_uuid()
```

Semantics:

- This function does not expect any input argument.
- **return type:** string

Example 12-36 Autogenerate primary key values using UUID function

In certain applications like maintaining student records in a university, you can automatically generate the ID instead of providing the value. Consider a simple table schema with an id as the primary key and a name column.

```
CREATE TABLE myTable (id STRING AS UUID, name STRING, PRIMARY KEY (id))
```

Insert the following data into the table. You can use the `random_uuid` function to generate the primary key values.

```
INSERT INTO myTable values(random_uuid(),"Adam")
```

```
INSERT INTO myTable values(random_uuid(),"Lily")
```

Explanation:

The `random_uuid` function internally generates unique UUID during the INSERT operation and assigns them to the id fields. Run the query to select the elements from the table.

```
select * from myTable order by name
```

Output:

```
{"id":"ff7057c2-cda9-4f6b-b94f-227b259a94d3","name":"Adam"}  
{"id":"37166790-4470-4484-bfbb-66364e0ff807","name":"Lily"}
```

The output of the query displays the assigned UUID values against the student names.

The `random_uuid` function generates a random but unique 36-byte string. Consider the multi-region tables where the identity columns are unique to a single region. You can use the `random_uuid` function to generate a globally unique identity during the record insertion.

You can retrieve the UUID of the inserted record by using the RETURNING clause in the INSERT or UPDATE statement as follows:

```
INSERT INTO myTable values(random_uuid(),"Adam") RETURNING *
```

Output:

```
{"id":"9f05eb60-2fa7-4c32-a90a-64371961cb9d","name":"Adam"}
```

For more details on the RETURNING clause, see the [Upsert statement](#).

You can retrieve the generated value of an identity column using class methods from various Oracle NoSQL Database Drivers.

For example: If your application is using the Java SDK, you can use the `getGeneratedValue()` method, which returns the generated value if the operation creates a new value for an identity column. For more details, see `PutResult`.

If your application is using the Python SDK, you can use the `get_return_row()` method, which succeeds only if the row exists. For more details, see `PutRequest`.

Similarly, each language SDK exposes an interface for retrieving the generated value.

Functions on Rows

As described in the [Table Management](#) section, table rows are record values conforming to the table schema, but with some additional properties that are not part of the table schema. To extract the value of such properties, the functions listed in this section must be used.

Although the signature of these functions specifies AnyRecord as the type of the input parameter, the functions actually require a row as input. The only expression that returns a row is a row variable, that is, a table alias whose name starts with '\$'. The [Example: Updating TTL](#) section shows an example of using the `remaining_hours()` function, which is one of the row available functions.

modification_time function

The `modification_time` function allows you to see the most recent modification time (in UTC) of a row. The time is returned as a timestamp value of precision 3 (milliseconds). If the row has never been modified since its insertion, it returns the insertion time. You may find this useful in deployments where tables span multiple regions and the Oracle NoSQL Database cross-region agent is updating your table by replicating data from remote regions.

Syntax:

```
timestamp(3) modification_time (AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** timestamp (3)

Example 12-37 Fetch the last modified time details for a passenger record from the airline baggage tracking application table

In an airline application, you can schedule periodic back-ups of passenger travel records based on the last modification time. Only the records that have been modified after the last scheduled backup can be considered for the next iteration.

```
SELECT modification_time($u)
FROM BaggageInfo $u
WHERE ticketNo = 1762344493810
```

Explanation:

You can use the `modification_time` function to retrieve the last modified time details for a passenger record from the `BaggageInfo` table.

Output:

```
+-----+
|      Column_1      |
+-----+
| 2023-01-18T07:53:02.048Z |
+-----+
```


The query returns the information on the most recent modification time for the passenger with ticketNo **1762344493810** from the BaggageInfo table.

creation_time_millis function

Returns the creation time of the row in number of milliseconds since January 1, 1970 UTC.

Note

This feature is available starting with Release 26.1. To enable it, see Enabling new features.

Syntax:

```
long creation_time_millis(AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value
- **Return Type:** long

Example: Fetch the creation time of a passenger record from the airline baggage tracking application table

You can use the creation_time_millis function to fetch the time (in milliseconds) when a passenger record was created.

```
SELECT creation_time_millis($u) AS creationtime_ms FROM BaggageInfo $u WHERE  
ticketNo =  
1762344493810
```

Explanation:

The query fetches the creation time (in milliseconds) of the row with ticket number 1762344493810.

Output:

```
{"creationtime_ms":1750662332583}  
1 row returned
```

creation_time function

Returns the creation time of a row in UTC (Coordinated Universal Time). The time is returned as a timestamp value of precision 3 (milliseconds).

Note

This feature is available starting with Release 26.1. To enable it, see Enabling new features.

Syntax:

```
timestamp(3) creation_time (AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value
- **Return Type:** timestamp (3)

Example: Fetch the creation time of a passenger record from the airline baggage tracking application table

You can use the `creation_time` function to fetch the time (in timestamp value) when a passenger record was created.

```
SELECT creation_time($u) AS CREATION_TIME FROM BaggageInfo $u WHERE ticketNo  
= 1762344493810
```

Explanation:

The query fetches the creation time (timestamp value) of the row with ticket number 1762344493810.

Output:

```
{"CREATION_TIME": "2025-06-23T07:05:32.583Z"}  
1 row returned
```

remaining_hours function

Returns the number of full hours remaining until the row expires. If the row has no expiration time, it returns a negative number.

Syntax:

```
integer remaining_hours(AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** integer

remaining_days function

Returns the number of full days remaining until the row expires. If the row has no expiration time, it returns a negative number.

Syntax:

```
integer remaining_days(AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** integer

expiration_time function

Returns the expiration time of the row, as a timestamp value of precision zero. If the row has no expiration time, it returns a timestamp set on January 1, 1970 UTC.

Syntax:

```
timestamp(0) expiration_time(AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** timestamp(0)

expiration_time_millis function

Returns the expiration time of the row, as the number of milliseconds since January 1, 1970 UTC. If the row has no expiration time, it returns zero.

Syntax:

```
long expiration_time_millis(AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** long

Example 12-38 TTL-related functions

In an airline application, you may want to back up the passenger records before the data expiration time or extend the expiration date to retain the information in the table for a longer period.

```
SELECT
remaining_hours($u) AS hours,
remaining_days($u) AS days,
expiration_time($u) AS expirytime,
expiration_time_millis($u) AS expirytime_ms
FROM BaggageInfo $u
WHERE ticketNo = 1762344493810
```

Explanation:

You can use the TTL expiration functions to check the expiration details (if any) of the rows containing the passenger records in the BaggageInfo table.

Output:

```
{"hours":376,"days":15,"expirytime":"2023-02-04T00:00:00.000Z","expirytime_ms":1675468800000}
```

The query returns the TTL information on the row that contains the passenger data for ticketNo **1762344493810** from the BaggageInfo table.

Here, the row expires after 15 days. The same information is displayed in hours, timestamp value, and number of milliseconds since January 1, 1970, UTC using the row functions described above.

You can update the expiration day/time of a row using the UPDATE statement as follows:

```
UPDATE BaggageInfo $u
SET TTL remaining_days($u) + 15 days
WHERE ticketNo = 1762344493810
```

The above statement extends the life of the row by 15 days. In this example, the remaining days until the expiry of the row is 15 days, to which the above statement adds 15 more days, effectively rendering the row to expire after 30 days. You can use the `remaining_hours` function and add hours to it to extend the expiration of a row by hours.

For more details on updating the TTL values, see [Example: Updating TTL](#).

The following functions allow you to see how the data is distributed across the store and collect statistics:

version or row_version function

The version or `row_version` function allows you to see the version of a given row. Every time a row is inserted or modified, the row version value is automatically calculated and the previous row version value is replaced with the new value. The row version value is unique for every row within the table and cannot be modified manually. Both `version` and `row_version` functions fetch the same result and can be used interchangeably in a query. You can use the `version` and `row_version` functions in both SELECT and WHERE clauses as demonstrated in the **Example** section.

Syntax:

```
binary version(AnyRecord)
```

```
binary row_version(AnyRecord)
```

Semantics:

- **AnyRecord:** The version or `row_version` function expects a row as the input value.
- **return type:** The version or `row_version` function returns a binary value. To use the version as a string in a query, you must cast the value to a string using the [CAST expression](#).

Example 12-39 Fetch the version of the row with passenger ticket number 1762344493810

```
SELECT row_version($t) from Baggageinfo $t WHERE ticketNo =
1762344493810
```

Explanation:

In this query, you fetch the version of the row that includes passenger data with ticket number 1762344493810. You supply the row as an input argument to the `row_version` function. Using the version function instead will also return the same result.

Output:

```
{"Column_1":"r00ABXcsACEZ1TncQVVExr5UlPEJixAJAAAAAAAAAKDIBAwAAAAEAAAABAAAAAA01xE="}
```

In the following example, you use the UPDATE statement to modify the confNo field in a passenger's data and fetch the updated row version.

```
UPDATE baggageinfo $t
SET confNo = "LE6J6Z"
WHERE ticketNo = 1762344493810
RETURNING row_version($t) as rowVersion
```

Output:

```
{"rowVersion":"r00ABXcsACEZ1TncQVVExr5UlPEJixAJAAAAAAAAATP8BAwAAAAEAAAABAAAAAA
aylA="}
```

After the row is updated, the version of the corresponding row is replaced with a new value automatically by Oracle NoSQL Database server.

In the following example, you use the row_version function in a WHERE clause to fetch the passenger data associated with a row version.

```
SELECT ticketNo AS TICKET, fullName AS NAME
FROM baggageinfo $t
WHERE cast(row_version($t) as String)
="r00ABXcsACEZ1TncQVVExr5UlPEJixAJAAAAAAAAAKDIBAwAAAAEAAAABAAAAAA01xE="
```

Notice that you cast the result of the row_version function to a string using the CAST expression to compare the versions. In this example, a previous row version is used to fetch the passenger data. As the previous row version no longer exists when the data is updated, there is no passenger data associated with it.

0 row returned

You can also use the row_version function in other NoSQL expressions. The following example determines if passenger data has been modified using the row_version function in a [CASE expression](#):

```
SELECT ticketNo AS TICKET, modification_time($t) AS TIME,
CASE
WHEN CAST(row_version($t) as String) =
"r00ABXcsACEZ1TncQVVExr5UlPEJixAJAAAAAAAAAKDIBAwAAAAEAAAABAAAAAA01xE="
THEN "row is not modified"
ELSE "row is modified"
END AS REPORT
FROM baggageinfo $t
WHERE ticketNo = 1762344493810;
```

You first cast the result of the row_version function to a string using the CAST expression and then compare the value with the previously stored version. Depending on the result of the

string comparison operation, the corresponding message is displayed. This query also uses the [modification_time](#) function to see the most recent modification time of the passenger's data.

Output:

```
{"TICKET":1762344493810,"TIME":"2024-04-04T09:42:52.128Z","REPORT":"row is modified"}
```

last_write_metadata function

The `last_write_metadata` function retrieves the user-supplied metadata for a row. The metadata is created when you supply a JSON string as an input during insert, update, or delete operations using APIs. The `last_write_metadata` function then returns a JSON construct, allowing you to access elements within the metadata using JSON dot notation and JSON query operators.

Note

If a row doesn't contain metadata, the function returns a null.

Syntax:

```
json last_write_metadata(AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **returnType:** JSON

Example 12-40 Fetch the user-supplied metadata for a given row

You can use the `last_write_metadata()` to fetch the metadata for the desired row of a table if the metadata exists.

```
SELECT last_write_metadata($t) AS info FROM stream_acct $t WHERE acct_Id=1
```

Output:

```
{"info":{"modified_by":"John Doe","reviewed_in":"Q1","update_reason":"Account details updated"}}
```

For more examples, see [Using SQL commands on User-Supplied Metadata](#).

integer shard (AnyRecord)

The `shard` function allows you to retrieve the shard ID in which a given rows of data is stored. It returns an integer value. For more information on shard ID, see the [Viewing Key Distribution Statistics](#) section in the Administrator's Guide.

Syntax:

```
integer shard (AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** integer

Example 12-41 Determine the shard details where the passenger record is stored

You may identify potential storage hotspots or an imbalance in your Oracle NoSQL Database cluster using the function. For example, you may notice that a particular shard seems to consume more storage than any other shard.

```
SELECT shard($u) AS Shard
FROM BaggageInfo $u
WHERE ticketNo = 1762344493810
```

Explanation:

In an airline application, you use the shard function to determine the shard details where the passenger record is stored.

Output:

```
{"Shard":1}
```

The query returns the shard in which the row with ticketNo **1762344493810** is stored in the BaggageInfo table.

You can compare the number of records in that shard with the other shards in your store by issuing the following command.

```
SELECT count(ticketNo) AS Shard_count
FROM BaggageInfo $u
WHERE shard($u) =1
```

Output:

```
{"Shard_count":2}
```

partition function

All data in the KVStore is accessed by one or more Keys. A Key might be a column in a table, or it might be the key portion of a Key/Value pair. The Keys are placed in logical containers called partitions, and each shard contains one or more partitions. For more details, see Partitions in the *Concepts Guide*.

The partition function allows you to see the partition id in which a given rows of data is stored.

Syntax:

```
integer partition (AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.

- **return type:** integer

Example 12-42 Determine the partitions in which the rows are stored

Use the partition function to determine the partitions in which the rows are stored in Oracle NoSQL Database.

```
SELECT partition($u) AS partition
FROM BaggageInfo $u
WHERE ticketNo = 1762344493810
```

Explanation:

In this example, the partition function returns the partition details of where the passenger record is stored. The query returns the partition in which the row with ticketNo **1762344493810** is stored in the BaggageInfo table.

Output:

```
{"partition":80}
```

The following query returns the list of partitions in the shard with id 1.

```
SELECT partition ($u) AS partition_count
FROM BaggageInfo $u
WHERE shard($u)=1
```

Output:

```
{"partition_count":80}
{"partition_count":131}
```

row_storage_size function

The row_storage_size function allows you to see the persistent storage size (in bytes) used by the given rows of data. It returns an integer value.

Syntax:

```
integer row_storage_size (AnyRecord)
```

Semantics:

- **AnyRecord:** This function expects a row as the input value.
- **return type:** integer

Example 12-43 Fetch the record size for a given row

You can use the row_storage_size function to obtain the record size for a given row.

```
SELECT row_storage_size($u) AS storage_size
FROM BaggageInfo $u
WHERE ticketNo = 1762344493810
```


Explanation:

In an airline application, you can use the `row_storage_size` function to determine the storage size of the individual passenger records.

Output:

```
{"storage_size":1123}
```

The query returns the storage size of the row containing the passenger record with ticketNo **1762344493810** in the `BaggageInfo` table. The storage size is displayed in bytes.

index_storage_size function

The `index_storage_size` function allows you to see the persistent storage size (in bytes) used by the index for the given row(s) of data. It returns an integer value. This function takes two arguments. The first argument is the table reference in which the index is created. The second argument is the index name for which the storage size is required. The index name is case-insensitive.

Syntax:

```
integer index_storage_size (AnyRecord, String)
```

Semantics:

- **AnyRecord:** A table reference in which the index is created as its first argument.
- **string:** The name of the index.
- **return type:** integer

Example 12-44 Determine the storage size of the individual index on a table

You can use the `index_storage_size` function to determine the storage size of the individual index on a table. The storage size of each index must be queried separately. To determine the total storage size of all the indexes on a table, you have to call the function for every index.

```
SELECT index_storage_size($u,"idx_contact")
FROM baggageInfo $u
WHERE ticketNo = 1762344493810
```

Explanation:

In an airline application, you can create an index for required fields on your Oracle NoSQL Database table. You use the `index_storage_size` function to retrieve the storage size of each index.

In this example, an index is created on the `contactPhone` field in the table. The table name `BaggageInfo` table and the index name `idx_contact` are supplied as arguments to the function.

For more information on Indexes, see [About Indexes](#).

Output:

```
+-----+
| Column_1 |
```

```
+-----+
|      40 |
+-----+
```

The query returns the storage size of the row containing the passenger record with ticketNo **1762344493810** in the BaggageInfo table. The storage size is displayed in bytes.

Example 12-45 Fetch the total number of bytes used to store all the rows of the table in a partition

In an airline application, you can get a detailed view of how your storage is distributed for your table across the partitions in your cluster by using the following query. For each partition, the query retrieves the total number of bytes used to store all the rows of the table contained in that partition.

```
SELECT
partition($u) AS partition,
sum(row_storage_size($u)) AS sum
FROM BaggageInfo $u
GROUP BY partition($u)
```

Explanation:

In this example, the partition function determines the partitions in which the passenger records are stored in the BaggageInfo table. The sum function computes the storage size of each row and sums up the result.

Output:

```
+-----+-----+
| partition|    sum |
+-----+-----+
|      80|    1123|
|     131|    1115|
+-----+-----+
```

The first column in the output is the list of partition numbers, and the second column is the current size of those partitions.

Functions on GeoJson Data

The GeoJson specification defines the structure and content of JSON objects that are supposed to represent geographical shapes on earth (called geometries). The following functions interpret the JSON objects as geometries and allow the search for rows containing geometries that satisfy certain conditions.

For more information on the functions and examples, see Managing GeoJSON data in the *Developers Guide*.

boolean geo_intersect(any*, any*)

Raises an error if it can be detected at compile time that an operand will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns false if any operand returns 0 or more than 1 items.

- Returns NULL if any operand returns NULL.
- Returns false if any operand returns an item that is not a valid geometry object.

Finally, if both operands return a single geometry object, it returns true if the 2 geometries have any points in common; otherwise false.

boolean geo_inside(any*, any*)

Raises an error if it can be detected at compile time that an operand will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns false if any operand returns 0 or more than 1 items.
- Returns NULL if any operand returns NULL.
- Returns false if any operand returns an item that is not a valid geometry object (however, if it can be detected at compile time that an operand will not return a valid geometry, an error is raised).
- Returns false if the second operand returns a geometry object that is not a polygon.

Finally, if both operands return a single geometry object and the second geometry is a polygon, it returns true if the first geometry is completely contained inside the second polygon, i.e., all its points belong to the interior of the polygon; otherwise false. The interior of a polygon is all the points in the polygon area except the points on the linear rings that define the polygon's boundary.

boolean geo_within_distance(any*, any*, double)

Raises an error if it can be detected at compile time that any of the first two operands will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns false if any of the first two operands returns 0 or more than 1 items.
- Returns NULL if any of the first two operands returns NULL.
- Returns false if any of the first two operands returns an item that is not a valid geometry object.

Finally, if both of the first two operands return a single geometry object, it returns true if the first geometry is within a distance of N meters from the second geometry, where N is the number returned by the third operand; otherwise false. The distance between 2 geometries is defined as the minimum among the distances of any pair of points where the first point belongs to the first geometry and the second point to the second geometry. If N is a negative number, it is set to 0.

boolean geo_near(any*, any*, double)

geo_near is converted internally to geo_within_distance plus an (implicit) order-by the distance between the two geometries. However, if the query has an (explicit) order-by already, no ordering by distance is performed. The geo_near function can appear in the WHERE clause only, where it must be a top-level predicate, i.e, not nested under an OR or NOT operator.

double geo_distance(any*, any*)

Raises an error if it can be detected at compile time that an operand will not return a single valid geometry object. Otherwise, the runtime behavior is as follows:

- Returns -1 if any operand returns zero or more than 1 items.
- Returns NULL if any operand returns NULL.
- Returns -1 if any of the operands is not a geometry.

Otherwise it returns the geodetic distance between the 2 input geometries. The returned distance is the minimum among the distances of any pair of points where the first point belongs to the first geometry and the second point to the second geometry. Between two such points, their distance is the length of the geodetic line that connects the points.

boolean geo_is_geometry(any*)

- Returns false if the operand returns zero or more than 1 items.
- Returns NULL if the operand returns NULL.
- Returns true if the input is a single valid geometry object. Otherwise, false.

Functions on Strings

This section describes various functions on strings.

To follow along with the examples in the section, create users table that provides information about users. See the [seq_distinct function](#) in *Functions on Sequences* section for users table creation.

The following string functions are supported:

Table 12-7 String Functions

Function	Description
substring Function	Extracts a substring from a string based on a specified starting position and length.
concat Function	Returns a single value of concatenating the two arguments.
upper Function	Converts the characters in the string to uppercase.
lower Function	Converts the characters in the string to lowercase.
trim Function	Returns a string with leading, trailing, or both sets of characters trimmed.
ltrim Function	Returns a string with leading characters trimmed.
rtrim Function	Returns a string with trailing characters trimmed.
length Function	Returns the length of the character string.
contains Function	Returns a boolean indicating if the search string is present in the source string.
starts_with Function	Returns a boolean indicating if the source string begins with the search string.
ends_with Function	Returns a boolean indicating if the source string ends with the search string.
index_of Function	Returns the position of the first character of the search string at its first occurrence.
replace Function	Returns the source string with each occurrence of the search string replaced by the replacement string.
reverse Function	Returns the characters of the source string in the reverse order.

substring Function

The substring function extracts a string from a given string according to a given numeric starting position and a given numeric substring length.

Syntax

```
returnvalue substring (source, position [, substring_length] )
```

```
source ::= any*
position ::= integer*
substring_length ::= integer*
returnvalue ::= string
```

Semantics

source

The input string from which the substring should be extracted. This argument is implicitly cast to a sequence of strings.

position

This argument indicates the starting point of the substring within the source. The first character of the source string has position 0.

An error is thrown if a non-integer value is supplied for the position.

substring_length

This argument indicates the length of the substring starting from the position value. If the supplied value is greater than the length of the source, then the length of the source is assumed for this argument.

An error is thrown if a non-integer value is supplied for the substring_length.

returnvalue

Returns an empty string ("") if the function did not return any characters.

Returns an empty string ("") if the substring_length is less than 1.

Returns NULL if the source argument is NULL.

Returns NULL if the position argument is less than 0 or greater or equal to the source length.

Example 12-46 substring Function

In this example, the first character in the firstname is selected from the users table. Notice that to select the first character, we have provided the value 0 for the position argument.

```
SELECT substring(firstname,0,1) as Initials FROM users
```

Output

```
+-----+
| Initials |
+-----+
| J         |
| P         |
| M         |
+-----+
```

Example 12-47 substring Function

This example illustrates that providing a negative value for the position argument will result in a NULL output value.

```
SELECT substring (firstname, -5, 4) FROM users
```

Output

Column_1
NULL
NULL
NULL

Example 12-48 substring Function

In this example, we select the first 4 characters from the firstname in the users table.

```
SELECT substring (firstname, 0, 4) FROM users
```

Output

Column_1
John
Pete
Mary

Example 12-49 substring Function

In this example, we select 100 characters starting from position 2. Notice that even though none of the rows has more than 5 characters in firstname, still we get the output up to the length of the source starting from position 2.

```
SELECT substring (firstname, 2, 100) FROM users
```

Output

Column_1
hn
ter
ry

Example 12-50 substring Function

In this example, the `substring_length` argument is not provided as it is optional. In such cases, we get the complete substring starting from the given position.

```
SELECT substring (firstname, 2) FROM users
```

Output

```
+-----+  
| Column_1 |  
+-----+  
| hn      |  
| ter     |  
| ry      |  
+-----+
```

concat Function

The `concat` function returns `arg1` concatenated with `arg2`. Both `arg1` and `arg2` can be of any data type.

Syntax

```
returnvalue concat (source,[source*])
```

`source ::= any*`

`returnvalue ::= boolean`

Semantics**source**

The input values that are joined to get a character string. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns the character string made by joining its character string operands in the order given. If any of the arguments is a sequence, then all the items are concatenated to the result in the order they appear in the sequence.

If all the arguments are empty sequence, then an empty sequence is returned.

If all the arguments are NULL, then a NULL is returned. This is because a NULL argument is converted to an empty string during concatenation unless all arguments are NULL, in which case the result is NULL. So NULL can result only from the concatenation of two or more NULL values.

Note

For security/denial of service reasons the maximum number of chars of the returned string will be less than `STRING_MAX_SIZE = 2^18 - 1` in chars i.e. 512kb. If the number of chars exceeds this number, then a runtime query exception is thrown.

Example 12-51 concat function

This example joins id, firstname, and lastname into a single string and provides the output. Notice that id, which is an integer type, also gets concatenated with the string values.

```
SELECT concat(id, firstname, lastname) AS name FROM users
```

Output

```
+-----+  
|   name   |  
+-----+  
| 10JohnSmith |  
| 30PeterPaul  |  
| 20MaryAnn   |  
+-----+
```

upper Function

The upper function converts all the characters in a string to uppercase.

Syntax

```
returnvalue upper (source)
```

source ::= any*

returnvalue ::= string

Semantics**source**

The input string that should be converted to uppercase. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns NULL if the source argument is NULL.

Returns NULL if the source argument is an empty sequence or a sequence with more than one item.

Note

If you want to convert a sequence with more than one item, see the Sequence Transform Expressions section.

Example 12-52 upper Function

In this example, the lastname field is converted to uppercase.

```
SELECT id, firstname, upper(lastname) FROM users
```


Output

```

+----+-----+-----+
| id | firstname | Column_3 |
+----+-----+-----+
| 10 | John      | SMITH     |
| 20 | Mary      | ANN       |
| 30 | Peter     | PAUL      |
+----+-----+-----+

```

lower Function

The lower function converts all the characters in a string to lowercase.

Syntax

```
returnvalue lower (source)
```

```
source ::= any*
```

```
returnvalue ::= string
```

Semantics**source**

The input string that should be converted to lowercase. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns NULL if the source argument is NULL.

Returns NULL if the source argument is an empty sequence or a sequence with more than one item.

Note

If you want to convert a sequence with more than one item, see the Sequence Transform Expressions section.

Example 12-53 lower Function

In this example, the lastname field is converted to lowercase.

```
SELECT id, firstname, lower(lastname) FROM users
```

Output

```

+----+-----+-----+
| id | firstname | Column_3 |
+----+-----+-----+
| 10 | John      | smith     |
| 20 | Mary      | ann       |
| 30 | Peter     | paul      |
+----+-----+-----+

```

trim Function

The trim function enables you to trim leading or trailing characters (or both) from a string.

Syntax

```
returnvalue trim(source [, position [, trim_character]])
```

```
source ::= any*  
position ::= "leading"|"trailing"|"both"  
trim_character ::= string*  
returnvalue ::= string
```

Semantics

source

The input string that should be trimmed. This argument is implicitly cast to a sequence of strings.

If you provide only the source argument, then the leading and trailing blank spaces are removed.

position

This argument indicates whether leading or trailing or both leading and trailing characters should be removed. The following are the valid values that can be specified for this argument.

- If *leading* is specified, then the characters equal to the trim_character argument are removed from the beginning of the string.
- If *trailing* is specified, then the characters equal to the trim_character argument are removed at the end of the string.
- If *both* is specified, then the characters equal to the trim_character argument are removed from both the beginning and end of the string.
- If no value is specified, then *both* value is assumed.
- If any value other than the above valid values are specified, then NULL is returned.

trim_character

This argument specifies the characters that should be removed from the source string. If you do not specify this argument, then a blank space is taken as the default value.

Only one character is allowed for this argument. If there are more than one character, then the first character will be used.

If an empty string is specified, then no trimming happens.

return_value

Returns NULL if any of the arguments is NULL.

Returns NULL if any argument is an empty sequence or a sequence with more than one item.

Example 12-54 trim function

Create this table and insert values in it to run the trim, ltrim, and rtrim function examples.

```
CREATE TABLE trim_demo (  
  id INTEGER,  
  name STRING,  
  yearofbirth STRING,
```

```
PRIMARY KEY (id)
)
```

```
INSERT INTO trim_demo VALUES (10, " Peter ", 1980)
```

```
INSERT INTO trim_demo VALUES (20, "Mary", 1973)
```

```
INSERT INTO trim_demo VALUES (30, " Oliver", 2000)
```

```
INSERT INTO trim_demo VALUES (40, "John ", 2000)
```

```
SELECT * FROM trim_demo
```

Output

id	name	yearofbirth
10	Peter	1980
20	Mary	1973
30	Oliver	2000
40	John	2000

Example 12-55 trim Function

In this example, the id and yearofbirth are selected from the trim_demo table. Notice that the zeros at the end of the yearofbirth are removed using the trim function.

```
SELECT id, trim(yearofbirth,"trailing",'0') FROM trim_demo
```

Output

id	Column_2
10	198
20	1973
30	2
40	2

Example 12-56 trim Function

In this example, '19' is provided as the trim_character. However, as per semantics, only the first character '1' will be considered for trimming.

```
SELECT id, trim(yearofbirth,"leading",'19') FROM trim_demo
```

Output

```

+----+-----+
| id | Column_2 |
+----+-----+
| 10 | 980      |
| 20 | 973      |
| 30 | 2000     |
| 40 | 2000     |
+----+-----+

```

ltrim Function

The ltrim function enables you to trim leading characters from a string.

Syntax

```
returnvalue ltrim(source)
```

source ::= any*

returnvalue ::= string

Semantics**source**

The input string that should be trimmed. The leading spaces in this string are removed. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns NULL if the source argument is NULL.

Returns NULL if the source argument is an empty sequence or a sequence with more than one item.

Example 12-57 trim Function

This example demonstrates ltrim function. Notice that the empty spaces at the beginning are removed but the empty spaces at the end are not removed.

Note

You can use JSON query output mode so that the empty spaces are visible.

```
MODE JSON
```

```
SELECT id, ltrim(name) FROM trim_demo
```

Output

```

{"id":10,"Column_2":"Peter "}
{"id":20,"Column_2":"Mary"}

```

```
{"id":30,"Column_2":"Oliver"}  
{"id":40,"Column_2":"John  "}
```

rtrim Function

The `rtrim` function enables you to trim trailing characters from a string.

Syntax

```
returnvalue rtrim(source)
```

```
source ::= any*  
returnvalue ::= string
```

Semantics

source

The input string that should be trimmed. The trailing spaces in this string are removed. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns NULL if the source argument is NULL.

Returns NULL if the source argument is an empty sequence or a sequence with more than one item.

Example 12-58 trim Function

This example demonstrates `rtrim` function. Notice that the empty spaces at the end are removed but the empty spaces at the beginning are not removed.

Note

You can use JSON query output mode so that the empty spaces are visible.

```
MODE JSON
```

```
SELECT id, rtrim(name) FROM trim_demo
```

Output

```
{"id":10,"Column_2":"  Peter"}  
{"id":20,"Column_2":"Mary"}  
{"id":30,"Column_2":"  Oliver"}  
{"id":40,"Column_2":"John"}
```

length Function

The `length` function returns the length of a character string. The `length` function calculates the length using the UTF character set.

Syntax

returnvalue **length**(source)

source ::= any*

returnvalue ::= integer

Semantics

source

The input string for which the length should be determined. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns NULL if the source argument is NULL.

Returns NULL if the source argument is an empty sequence or a sequence with more than one item.

Note

Characters that are represented on 32 or more bits, the length is considered 1, while Java String.length() returns 2 for UTF32 chars, 4 for UTF64, etc.

Example 12-59 length Function

In this example, the length of the first name is selected from the users table.

```
SELECT firstname, length(firstname) as length FROM users
```

Output

firstname	length
John	4
Mary	4
Peter	5

contains Function

The contains function indicates whether or not a search string is present inside the source string.

The contains function is a Boolean function that performs a simple, literal substring search. The function checks the source string character by character to determine whether the exact sequence of characters from the search string appears anywhere. All characters in the search string are treated as plain text. The contains function does not use regular expressions or SQL wildcards, such as % or _.

Syntax

```
returnvalue contains(source, search_string)
```

```
source ::= any*
search_string ::= any*
returnvalue ::= boolean
```

Semantics**source_string**

The input string to be searched. This argument is implicitly cast to a sequence of strings.

search_string

The string that should be searched in the source. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns true if the search_string is found inside the source string.

Returns false if the search string is not found, or if any argument is an empty sequence or a sequence with more than one item.

Returns NULL if the source or search_string argument is NULL.

Example 12-60 contains Function

In this example, the firstname field values that contain the string "ar" in it is indicated as true.

```
SELECT firstname, contains(firstname,"ar") FROM users
```

Output

```
+-----+-----+
| firstname | Column_2 |
+-----+-----+
| John      | false    |
| Peter     | false    |
| Mary      | true     |
+-----+-----+
```

starts_with Function

The starts_with function indicates whether or not the source string begins with the search string.

Syntax

```
returnvalue starts_with(source, search_string)
```

```
source ::= any*
search_string ::= any*
returnvalue ::= boolean
```

Semantics

source

The input string to be searched. This argument is implicitly cast to a sequence of strings.

search_string

The string that should be searched in the source. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns true if source begins with search_string else returns false.

Returns false if any argument is an empty sequence or a sequence with more than one item.

Returns NULL if source or search_string is NULL.

Example 12-61 starts_with Function

In this example, the firstname field values that starts with the string "Pe" is indicated as true.

```
SELECT firstname, starts_with(firstname,"Pe") FROM users
```

Output

firstname	Column_2
John	false
Peter	true
Mary	false

ends_with Function

The ends_with function indicates whether or not the source string ends with the search string.

Syntax

```
returnvalue ends_with(source, search_string)
```

source ::= any*

search_string ::= any*

returnvalue ::= boolean

Semantics

source

The input string to be searched. This argument is implicitly cast to a sequence of strings.

search_string

The string that should be searched in the source. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns true if source ends with search_string else returns false.

Returns false if any argument is an empty sequence or a sequence with more than one item.
Returns NULL if source or search_string is NULL.

Example 12-62 ends_with Function

In this example, the firstname field values that ends with the string "hn" is indicated as true.

```
SELECT firstname, ends_with(firstname,"hn") FROM users
```

Output

firstname	Column_2
John	true
Peter	false
Mary	false

index_of Function

The index_of function determines the position of the first character of the search string at its first occurrence, if any.

Syntax

```
returnvalue index_of(source, search_string [, start_position])
```

```
source ::= any*
search_string ::= any*
start_position ::= integer*
returnvalue ::= integer
```

Semantics

source

The input string to be searched. This argument is implicitly cast to a sequence of strings.

search_string

The string that should be searched in the source. This argument is implicitly cast to a sequence of strings.

start_position

An optional integer indicating, numerically, the position in the source from where the search should begin.

The default start_position is 0 which is also the position of the first character in the source.
If a negative value is supplied to start_position then 0 is assumed.

returnvalue

Returns the position of the first character of the search string at its first occurrence.

Returns -1 if search_string is not present in source.

Returns 0 for any value of source if the search_string is of length 0.

Returns NULL if any argument is NULL.

Returns NULL if any argument is an empty sequence or a sequence with more than one item.

Returns error if start_position argument is not an integer.

Note

The returnvalue is relative to the beginning of source, regardless of the value of start_position.

Example 12-63 index_of Function

In this example, the index of "r" is selected in the firstname.

In the output, John has no occurrence of "r" so -1 is returned. Peter and Mary has "r" at 4 and 2 position respectively.

```
SELECT firstname, index_of(firstname,"r") FROM users
```

Output

firstname	Column_2
John	-1
Peter	4
Mary	2

Example 12-64 index_of Function

In this example, the index of "e" is selected in the firstname. In the output, notice that although "e" occurs twice in Peter, only the position of the first occurrence is returned.

```
SELECT firstname, index_of(firstname,"e") FROM users
```

Output

firstname	Column_2
John	-1
Peter	1
Mary	-1

replace Function

The replace function returns the source with every occurrence of the search string replaced with the replacement string.

Syntax

```
returnvalue replace(source, search_string [, replacement_string])
```

```

source ::= any*
search_string ::= any*
replacement_string ::= any*
returnvalue ::= string

```

Semantics

source

The input string that should be searched. This argument is implicitly cast to a sequence of strings.

search_string

The string that should be searched in the source. This argument is implicitly cast to a sequence of strings.

replacement_string

The string that should be substituted in place of search_string in the source. This is an optional argument. If replacement_string is omitted or empty sequence, then all occurrences of search_string are removed from source. The result will be checked so that the result would not be bigger than STRING_MAX_SIZE = 2¹⁸ - 1 in chars ie. 512kb, if that is the case a runtime query exception is thrown. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns source if the search_string argument is NULL.

Returns NULL if source argument is NULL.

Returns NULL if either source or search_string argument is an empty sequence.

Returns NULL if any argument is a sequence with more than one item.

Example 12-65 replace Function

In this example, the string "e" is replaced with "X" in all the occurrences in firstname. Notice the occurrence of "X" in Peter.

```
SELECT firstname, replace(firstname,"e","X") FROM users
```

Output

```

+-----+-----+
| firstname | Column_2 |
+-----+-----+
| John      | John      |
| Peter     | PXtXr     |
| Mary      | Mary      |
+-----+-----+

```

Example 12-66 replace Function

In this example, the string "ar" is replaced with "urph". Notice that in the source the remaining characters after the search_string are retained for output. This yields the output for "Mary" as "Murphy".

```
SELECT firstname, replace(firstname,"ar","urph") FROM users
```

Output

```

+-----+-----+
| firstname | Column_2 |
+-----+-----+
| John      | John      |
| Peter     | Peter     |
| Mary      | Murphy    |
+-----+-----+

```

Example 12-67 replace Function

In this example, the replacement_string is not specified. Since the replacement_string is not specified, the search_string is removed and the remaining source is displayed.

```
SELECT firstname, replace(firstname,"oh") FROM users
```

Output

```

+-----+-----+
| firstname | Column_2 |
+-----+-----+
| John      | Jn        |
| Peter     | Peter     |
| Mary      | Mary      |
+-----+-----+

```

reverse Function

The reverse function returns the characters of the source string in reverse order, where the string is written beginning with the last character first. For example, the reverse order for "welcome" is "emoclew".

Syntax

```
returnvalue reverse(source)
```

```
source ::= any*
```

```
returnvalue ::= string
```

Semantics**source**

The input string for which the characters should be reversed. This argument is implicitly cast to a sequence of strings.

returnvalue

Returns NULL if the source argument is NULL.

Returns NULL if the source argument is an empty sequence or a sequence with more than one item.

Example 12-68 reverse Function

In this example, the first name is displayed along with its reverse order.

```
SELECT firstname, reverse(firstname) FROM users
```

Output

firstname	Column_2
John	nhoJ
Peter	reteP
Mary	yraM

Function to Convert String to JSON

parse_json function

The function `parse_json` converts a string argument to a JSON instance. The input string must be a comma-separated list of one or more name-value pairs.

Syntax:

```
json parse_json(string)
```

Semantics:

The input string argument must be a valid JSON text. The `parse_json` function parses the string and converts it to a JSON object.

For more details on JSON data type in Oracle NoSQL Database data model, see [JSON Data Type](#). An error is displayed if an incorrect JSON text is supplied in the string argument (for example, a missing colon to separate the name and field values).

Example 12-69 Consider a user data table for a library application. Currently, the subscription details are in a JSON document, which is stored as a string. You want to add them as a JSON object

To achieve this, consider the following schema for the table:

```
CREATE TABLE userslib (id LONG GENERATED BY DEFAULT AS IDENTITY, details
JSON, Book1 json,
      Book2 Json, Book3 Json, PRIMARY KEY (id))
```

You can declare the subscription ID as an IDENTITY column. The current subscription details can be included as a JSON object in the `details` field. Assuming, three books are allowed per subscription at any point, you can update the details of the borrowed books as JSON objects in the corresponding *Book* fields.

You use the `parse_json` function to convert the subscription data of a user, which is in a string format to a JSON object and insert it into the JSON field as follows:

```
mode json -pretty
```

```
insert into userslib (details) values
(parse_json("{\"firstName\":\"John\",\"lastName\":\"Smith\",\"DOB\":\"22-2-1995\",
\"address\":{\"city\":\"Santa Cruz\",
\"number\":101,\"contactphone\":\"408-453-8955\",\"state\":\"CA\",\"street\":\"Pacific Ave\",
\"zip\":95008},\"email\":\"john.smith@reachmail.com\"}")) RETURNING *
```

Explanation:

You must provide a string that is a valid JSON text. In this example, the INSERT statement parses the string using the `parse_json` function to create JSON objects, which are then updated as elements in the details field. Notice that the value for the `id` field, which is the primary key is auto-generated as you have declared it as an IDENTITY column. For more details, see [Inserting Rows with an IDENTITY Column](#). Also, since you have not provided any values for the books fields in this example, they are populated with NULL values.

Output:

```
{
  "id" : 2,
  "details" : {
    "DOB" : "22-2-1995",
    "address" : {
      "city" : "Santa Cruz",
      "contactphone" : "408-453-8955",
      "number" : 101,
      "state" : "CA",
      "street" : "Pacific Ave",
      "zip" : 95008
    },
    "email" : "john.smith@reachmail.com",
    "firstName" : "John",
    "lastName" : "Smith"
  },
  "Book1" : null,
  "Book2" : null,
  "Book3" : null
}
```

Example 12-70 Fetch from the library application, the book titles of the borrowed books for a user

In the `userslib` table above, update the book fields for a user with the details of the books that are borrowed from the library.

```
UPSERT into userslib values (2, {"DOB":"22-2-1995","address":{"city":"Santa Cruz",
"contactphone":"408-453-8955","number":101,"state":"CA","street":"Pacific Ave",
"zip":95008},"email":"john.smith@reachmail.com","firstName":"John","lastN
```

```
ame": "Smith"}, '{"doc":{"title":"A Tale of two cities", "author":"Charles Dickens", "site":"brooks.publishers.com"}}', '{"doc":{"title":"Harry Potter", "author":"J K Rowling", "site":"brooks.publishers.com"}}', '{"doc":{"title":"Percy Jackson", "author":"Rick Riordan", "site":"brooks.publishers.com"}}' RETURNING *
```

If any JSON data is in a string format inadvertently while updating the book details, the operation still succeeds as the string data is a valid JSON. However, the fields are populated with the unparsed string as follows:

```
{
  "id" : 2,
  "details" : {
    "DOB" : "22-2-1995",
    "address" : {
      "city" : "Santa Cruz",
      "contactphone" : "408-453-8955",
      "number" : 101,
      "state" : "CA",
      "street" : "Pacific Ave",
      "zip" : 95008
    },
    "email" : "john.smith@reachmail.com",
    "firstName" : "John",
    "lastName" : "Smith"
  },
  "Book1" : "{\"doc\":{\"title\":\"A Tale of two cities\", \"author\":\"Charles Dickens\", \"site\":\"brooks.publishers.com\"}}",
  "Book2" : {
    "doc" : "{\"title\":\"Harry Potter\", \"author\":\"J K Rowling\", \"site\":\"brooks.publishers.com\"}"
  },
  "Book3" : {
    "doc" : {
      "author" : "Rick Riordan",
      "site" : "brooks.publishers.com",
      "title" : "Percy Jackson"
    }
  }
}
```

Here, the Book1 field is populated as a complete string. The Book2 field has a doc attribute, which is JSON, however, the value is a string. The Book3 field is a proper JSON document.

You can use the `parse_json` function to select the JSON values from the table as follows:

```
SELECT
u.id, u.details.email,
parse_json(u.Book1).doc.title as title1,
parse_json(u.Book2.doc).title as title2,
u.Book3.doc.title as title3
FROM userslib u WHERE id=2
```

Explanation:

You can use the `parse_json` function to parse the string values in the JSON field. In the `Book1` field above, the value is a JSON document stored as a string. You parse the complete string to convert it to a JSON object and then select the `title` field. In the `Book2` field, the value of the `doc` attribute is a JSON document stored as a string. Here, you parse the attribute value to convert it to a JSON object and then select the `title` field. The `Book3` field is a valid JSON object, from which you can extract the `title` value directly using the [path expression](#).

Output:

```
{
  "id" : 2,
  "email" : "john.smith@reachmail.com",
  "title1" : "A Tale of two cities",
  "title2" : "Harry Potter",
  "title3" : "Percy Jackson"
}
```

Functions of Mathematical Operations

Learn about the Mathematical functions supported in Oracle NoSQL Database.

The mathematical functions are used to perform mathematical calculations on the input arguments. Here, the supplied arguments are expressions that resolve to numbers.

You can invoke the mathematical functions from the `SELECT/WHERE` clauses, and also from other Oracle NoSQL statements where function calls are allowed in the syntax. For example, you can supply mathematical functions as arguments to [Using Aggregate Functions](#).

The following functions are supported in the Oracle NoSQL Database:

Table 12-8 Mathematical functions

Function	Description
abs(n)	Returns the absolute value of <i>n</i> .
acos(n)	Returns the arc cosine of <i>n</i> expressed in radians.
asin(n)	Returns the arc sine of <i>n</i> expressed in radians.
atan(n)	Returns the arc tangent of <i>n</i> expressed in radians.
atan2(n1,n2)	Returns the arc tangent of arguments <i>n1</i> and <i>n2</i> expressed in radians.
ceil(n)	Returns the smallest integer that is greater than or equal to <i>n</i> .
cos(n)	Returns the cosine of an angle <i>n</i> specified in radians.
cot(n)	Returns the cotangent of an angle <i>n</i> specified in radians.
degrees(n)	Converts <i>n</i> from radians to degrees.
e()	Returns the value of the Euler number <i>e</i> .
exp(n)	Returns the exponential value of the expression <i>n</i> .
floor(n)	Returns the largest integer that is less than or equal to <i>n</i> .
ln(n)	Returns the natural logarithmic (base <i>e</i>) value of <i>n</i> .
log(n,b)	Returns the logarithmic base <i>b</i> value of <i>n</i> .
log10(n)	Returns the logarithmic base 10 value of <i>n</i> .

Table 12-8 (Cont.) Mathematical functions

Function	Description
pi()	Returns the value of <i>pi</i> .
power(n2,n1)	Returns the value of <i>n2</i> raised to the power <i>n1</i> .
radians(n)	Converts <i>n</i> from degrees to radians.
rand()	Returns a pseudo-random number between zero and one.
round(n [,d])	Rounds <i>n</i> to <i>d</i> places to the right of the decimal point.
sign(n)	Returns the sign of <i>n</i> .
sin(n)	Returns the sine of an angle <i>n</i> specified in radians.
sqrt(n)	Returns the square root of <i>n</i> .
tan(n)	Returns the tangent of an angle <i>n</i> specified in radians.
trunc(n [,d])	Returns the value of <i>n</i> truncated to <i>d</i> decimal places.

Note

- The mathematical function names are case-sensitive.
- All the mathematical functions expect numeric types (integer, long, float, double, and number) as arguments. These functions return NULL in the following cases. Here, the NULL refers to SQL NULL unless specified otherwise.
 - Any of the arguments resolves to NULL.
 - Any of the arguments is a field in a JSON column whose value is NULL, that is, JSON NULL.
 - Any of the arguments is a field in a JSON column, which does not exist.
 - Any of the arguments is a field in JSON collection, which does not exist.
 - Any of the arguments does not resolve to a numeric type.
- While using trigonometric functions such as [sin\(n\)](#), [cos\(n\)](#), and so forth, you must specify the input value in radians. You can convert an angle from degrees to radians using the [radians\(n\)](#) function. Similarly, for inverse trigonometric functions like [asin\(n\)](#), [acos\(n\)](#), and so forth, you can convert the resultant radians value to degrees using the [degrees\(n\)](#) function.

To follow along with the examples in the sections, you can create the tables and load data as described in the [Additional Examples](#) section.

abs function

The abs function returns the absolute value of the input expression.

Syntax:

```
number abs(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** number.
 - Returns the same type as the argument *n*.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns a positive infinity if *n* resolves to a positive/negative infinity.
 - Returns NaN if *n* resolves to NaN.
 - Returns a zero if *n* resolves to a positive/negative zero.
 - Returns Integer.MIN_VALUE if *n* resolves to an integer type value and is equal to Integer.MIN_VALUE.
 - Returns Long.MIN_VALUE if *n* resolves to a long type value and is equal to Long.MIN_VALUE.

Example 12-71 abs function

```
SELECT abs(n) FROM Userstocks
```

Value of n	Output
10	10
-10	10
1.67	1.67
-1.67	1.67

arc cosine function

The acos function is an inverse trigonometric function. The acos function returns the arc cosine (inverse of cosine) of the input expression.

Syntax:

```
double acos(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.

The resultant angle is expressed in radians and is in the range 0.0 through pi with the following exceptions:

 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN in the following cases:

- * n resolves to NaN
- * absolute value of n is greater than one
- * n resolves to positive/negative infinity

Example 12-72 arc cosine function

```
SELECT acos(n) FROM Archery
```

Value of n	Output
1	0.0
0	1.5707963267948966
sqrt(3)/2	0.5235987755982989
0.5	1.0471975511965979
-1	3.141592653589793

arc sine function

The asin function is an inverse trigonometric function. The asin function returns the arc sine (inverse of sine) of the input expression.

Syntax:

```
double asin(n)
```

Semantics:

- **n:** The argument n is an expression that resolves to a number.
- **return type:** double.
The resultant angle is expressed in radians and is in the range $-pi/2$ through $pi/2$ with the following exceptions:
 - Returns a NULL value in the following cases:
 - * n resolves to a NULL value
 - * n does not resolve to a number
 - Returns NaN in the following cases:
 - * n resolves to NaN
 - * absolute value of n is greater than one
 - * n resolves to positive/negative infinity
 - Results in a zero with the same sign as n , when n is zero.

Example 12-73 arc sine function

```
SELECT asin(n) FROM Archery
```

Value of n	Output
1	1.5707963267948966

Value of n	Output
sqrt(3)/2	1.0471975511965976
0.5	0.5235987755982989
-1	-1.5707963267948966

arc tan function

The atan function is an inverse trigonometric function. The atan function returns the arc tangent (inverse of tangent) of the input expression.

Syntax:

```
double atan(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.
The resultant angle is expressed in radians and is in the range $-\pi/2$ through $\pi/2$ with the following exceptions:
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN.
 - Results in a zero with the same sign as *n*, when *n* is zero.

Example 12-74 arc tan function

```
SELECT atan(n) FROM Archery
```

Value of n	Output
1	0.7853981633974483
sqrt(3)	1.0471975511965976

arc tan2 function

The atan2 function is an inverse trigonometric function. The atan2 function is a two-argument arc tangent (inverse of tan) function, which returns the arc tangent of a coordinate or point in two-dimensional space (*n1*, *n2*) calculated as [atan\(n2/n1\)](#).

Syntax:

```
double atan2(n1,n2)
```

Semantics:

- **n1, n2:** The arguments *n1* and *n2* are expressions that resolve to numbers.

- **return type:** double.
The resultant angle is expressed in radians and is in the range $-pi$ through pi with the following exceptions:
 - Returns a NULL value in the following cases:
 - * If either of the arguments resolves to a NULL value
 - * If either of the arguments does not resolve to a number
 - Returns a NaN if either of the arguments resolves to NaN.

Example 12-75 arc tan2 function

```
SELECT atan2(n1,n2) FROM Archery
```

Value of n1	Value of n2	Output
1	1	0.7853981633974483
1	sqrt(3)	1.0471975511965976
2	0	1.5707963267948966

ceil function

The ceil function returns the smallest integer that is greater than or equal to the specified expression.

Syntax:

```
number ceil(n)
```

Semantics:

- **n:** The argument n is an expression that resolves to a number.
- **return type:**
 - Returns the same type as the argument n .
 - Returns a NULL value in the following cases:
 - * n resolves to a NULL value
 - * n does not resolve to a number
 - Returns NaN if n resolves to NaN.
 - Returns a positive/negative infinity if n resolves to the respective infinity.
 - Returns a positive/negative zero if n resolves to the respective zero value.

Example 12-76 ceil function

```
SELECT ceil(n) FROM Userstocks
```

Value of n	Output
1.34	2.0
-1.34	-1.0

Value of n	Output
pi()	4.0

cosine function

The cos function is a trigonometric function. The cos function returns the cosine of an angle specified in radians.

Syntax:

```
double cos(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number (an angle expressed in radians).
- **return type:** double.
The resultant value is in the range -1 through 1 with the following exceptions:
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN or infinity(positive/negative).

Example 12-77 cosine function

```
SELECT cos(n) FROM Archery
```

Value of n	Output
0	1.0
radians(90)	6.123233995736766E-17
pi()	-1.0

cotangent function

The cot function is a trigonometric function. The cot function returns the cotangent of an angle specified in radians.

Syntax:

```
double cot(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number (an angle expressed in radians).
- **return type:** double.
 - Returns a NULL value in the following cases:

- * n resolves to a NULL value
- * n does not resolve to a number
- Returns NaN if n resolves to NaN or infinity(positive/negative).

Example 12-78 cotangent function

```
SELECT cot(n) FROM Archery
```

Value of n	Output
0	Infinity
radians(90)	6.123233995736766E-17
pi()/4	1.0000000000000002

degrees function

The degrees function converts the specified expression from radians to degrees.

Syntax:

```
double degrees(n)
```

Semantics:

- **n:** The argument n is expressed in radians. Here, n is an expression that resolves to a number.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * n resolves to a NULL value
 - * n does not resolve to a number
 - Returns NaN if n resolves to NaN.
 - Returns respective infinity if n resolves to positive/negative infinity.

Note

The conversion from radians to degrees does not generally return an exact value. For example, consider the query:

```
SELECT degrees(atan(sqrt(3))) AS ATAN60 FROM Archery
```

The degrees to radians conversion results in: {"ATAN60":59.99999999999999}

Example 12-79 degrees function

```
SELECT degrees(n) FROM Archery
```

Value of n	Output
pi() /2	90
0	0.0

Euler function

The e function returns the value of Euler number e, that is, 2.718281828459045.

Syntax:

```
double e()
```

Semantics:

- The e function does not expect any input arguments.
- **return type:** double.

Example 12-80 Euler function

```
SELECT e() AS EULER FROM Archery
```

Output:

```
+-----+-----+
| EULER      | 2.718281828459045 |
+-----+-----+
```

exp function

The exp function returns the exponential value of the given expression, that is, e raised to the power of the specified expression. Here, e is the base of the natural logarithm and has the value 2.718281828459045.

Syntax:

```
double exp(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN.
 - Returns positive infinity if *n* resolves to positive infinity.
 - Returns positive zero if *n* resolves to negative infinity.

Example 12-81 exp function

```
SELECT exp(n) FROM Phtable
```

Value of n	Output
0	1.0
2	7.38905609893065
-2	0.1353352832366127

floor function

The floor function returns the largest integer that is less than or equal to the specified expression.

Syntax:

```
number floor(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:**
 - Returns the same type as the argument *n*.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN.
 - Returns a positive/negative infinity if *n* resolves to the respective infinity.
 - Returns a positive/negative zero if *n* resolves to the respective zero value.

Example 12-82 floor function

```
SELECT floor(n) FROM Userstocks
```

Value of n	Output
1.34	1.0
-1.34	-2.0
pi()	3.0

ln function

The ln function returns the natural logarithmic value (base e) of the specified expression.

Syntax:

```
double ln(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN in the following cases:
 - * *n* resolves to NaN
 - * *n* is less than zero
 - Returns positive infinity if *n* resolves to positive infinity.
 - Returns negative infinity if *n* is positive/negative zero.

Example 12-83 ln function

```
SELECT ln(n) FROM PHTable
```

Value of n	Output
1	0.0
10	2.302585092994046
1.0E-7	-16.11809565095832

log function

The log function returns the logarithmic value with the specified base for the given expression.

Syntax:

```
double log(n,b)
```

Semantics:

- **n, b:** The arguments *n* and *b* are expressions that resolve to numbers. The base *b* can resolve to any positive value except zero and one. *n* can resolve to any positive value.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * If either of the arguments resolves to a NULL value
 - * If either of the arguments does not resolve to a number
 - Returns positive infinity if either of the arguments resolves to positive infinity.

Example 12-84 log function

```
SELECT log(n,b) FROM PHTable
```

Value of n	Value of b	Output
2	2	1.0
81	3	4.0
255	5	3.442980622208573

log10 function

The log10 function returns the logarithmic value with base 10 for the specified expression.

Syntax:

```
double log10(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN in the following cases:
 - * *n* resolves to NaN
 - * *n* is less than zero
 - Returns positive infinity if *n* resolves to positive infinity.
 - Returns negative infinity if *n* is positive/negative zero.
 - Returns *n* if the value of *n* is equal to 10^n , where *n* is an integer.

Example 12-85 log10 function

```
SELECT log10(n) FROM PHTable
```

Value of n	Output
1	0.0
10	1.0
1.0E-7	-7.0

pi function

The pi function returns the value of pi, that is, 3.141592653589793.

Syntax:

```
double pi()
```

Semantics:

- The `pi` function does not expect any input arguments.
- **return type:** double.

Example 12-86 pi function

```
SELECT pi() AS PI FROM Archery
```

Output:

```
+-----+-----+
| PI          | 3.141592653589793 |
+-----+-----+
```

power function

The `pow` function returns the value of the first expression raised to the power of the second expression.

Syntax:

```
double pow(n2,n1)
```

Semantics:

- ***n1*, *n2*:** The arguments *n1* and *n2* are expressions that resolve to numbers.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * If either of the arguments resolves to a NULL value
 - * If either of the arguments does not resolve to a number
 - Returns positive infinity if *n2* resolves to positive/negative infinity.
 - Returns positive infinity if *n1* resolves to positive infinity and zero if *n1* resolves to negative infinity.

Example 12-87 power function

```
SELECT pow(n2,n1) FROM PHTable
```

Value of <i>n</i>	Value of <i>b</i>	Output
2	4	16.0
2	0	1.0
-0.5	4	0.0625
0.5	-4	16.0

radians function

The `radians` function converts the specified expression from degrees to radians.

Syntax:

```
double radians(n)
```

Semantics:

- **n:** The argument *n* is expressed in degrees. Here, *n* is an expression that resolves to a number.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN.
 - Returns respective infinity if *n* resolves to positive/negative infinity.

Note

The conversion from degrees to radians does not generally return an exact value. For example, consider the query:

```
SELECT cot(radians(90)) COT90 FROM Archery
```

The degrees to radians conversion results in: {"COT90":6.123233995736766E-17}

Example 12-88 radians function

```
SELECT radians(n) FROM Archery
```

Value of n	Output
180	3.141592653589793
0	0.0

random function

The rand function returns a positive pseudo-random number. The resultant value can be greater than or equal to zero and less than one.

Syntax:

```
double rand()
```

Semantics:

- The rand function does not expect any input arguments.
- The rand function is evaluated for each row if used in a WHERE clause.
- **return type:** double.

Example 12-89 random function

```
SELECT rand() AS RANDOM from Archery WHERE sim=1
```

Output:

```
+-----+-----+
| RANDOM | 0.891655403699787 |
+-----+-----+
```

round function

The round function rounds the value of an input expression to the specified decimal places to the right of the decimal point.

Syntax:

```
double round(n [, d ])
```

Semantics:

- **n, d:** The arguments *n* and *d* are expressions that resolve to numbers. The value of *d* must be in the range of $-30 \leq d \leq 30$. Any value of *d* exceeding the maximum absolute value of 30 is truncated to 30 or -30 respectively. If *d* is not an integer, it is cast to an integer value. For example, round(123.456, 1.5) is considered round(123.456, 1), and round(123.456, 50) is treated as round(123.456, 30).
- **return type:** double.
Depending on the values of *n* and *d*, the round function behaves as follows:
 - If *d* is a positive value, the function rounds *n* to *d* places on the right of the decimal point. If the fractional part of *n* is 0.5 or greater, the value is rounded away from 0. For example, round(0.5) returns 1.0 and round(-0.5) returns -1.0.
 - If *d* is greater than or equal the number of digits to the right of the decimal point in *n*, the value of *n* is returned.
 - If *d* is not specified or zero, the function rounds *n* to zero decimal places.
 - If *d* is a negative value, the function rounds *n* to the left of the decimal point. If the number of digits in *n* to the left of the decimal point is less than or equal to the absolute value of *d*, then the function returns 0.0.
 - Returns zero if *n* is zero, regardless of the value of *d* (except when *d* is NULL).
 - Returns a NULL value in the following cases:
 - * If either of the arguments is a NULL value
 - * If either of the arguments does not resolve to a number
 - Returns *n* if either of the arguments resolves to a NaN or infinity.

Example 12-90 round function

```
SELECT round(n,d) FROM Userstocks
```

Value of n	Value of d	Output
100.331	2	100.33
100.367	2	100.37
111.567	0.5	112.0
111.567	4	111.567
10.361	0 or not specified	10.0
111.331	-2	100.0
111.331	-4	0.0
-100.331	2	-100.33
-111.331	-2	-100.0

sign function

The sign function returns the sign of the specified expression, that is, zero if the argument is zero, 1 if the argument is greater than zero, and -1 if the argument is less than zero.

Syntax:

```
double sign(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.
 - Returns 1.0 if *n* resolves to a positive number.
 - Returns -1.0 if *n* resolves to a negative number.
 - Returns 0.0 if *n* resolves to zero.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN.

Example 12-91 sign function

```
SELECT sign(n) FROM PHTable
```

Value of n	Output
1.89	1.0
-1.89	-1.0
0	0.0
sin(3*pi() 2)	-1.0

sine function

The `sin` function is a trigonometric function. The `sin` function returns the sine of an angle specified in radians.

Syntax:

```
double sin(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number (an angle expressed in radians).
- **return type:** double.
The resultant value is in the range -1 through 1 with the following exceptions:
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN if *n* resolves to NaN or infinity(positive/negative).
 - Returns a zero if *n* resolves to a positive/negative zero.

Example 12-92 sine function

```
SELECT sin(n) FROM Archery
```

Value of n	Output
radians(90)	1.0
pi()/4	0.7071067811865475
3*pi()/2	-1.0

square root function

The `sqrt` function returns the square root of the specified expression.

Syntax:

```
double sqrt(n)
```

Semantics:

- **n:** The argument *n* is an expression that resolves to a number.
- **return type:** double.
 - Returns a NULL value in the following cases:
 - * *n* resolves to a NULL value
 - * *n* does not resolve to a number
 - Returns NaN in the following cases:

- * n resolves to NaN
- * n is less than zero
- Returns positive infinity if n resolves to positive infinity.
- Returns the same value as the argument, if n resolves to positive/negative zero.

Example 12-93 square root function

```
SELECT sqrt(n) FROM Archery
```

Value of n	Output
16	4.0
12.57/ pi()	2.00028879648171

tangent function

The tan function is a trigonometric function. The tan function returns the tangent of an angle specified in radians.

Syntax:

```
double tan(n)
```

Semantics:

- **n:** The argument n is an expression that resolves to a number (an angle expressed in radians).
- **return type:** double.
The resultant value is in the range -1 through 1 with the following exceptions:
 - Returns a NULL value in the following cases:
 - * n resolves to a NULL value
 - * n does not resolve to a number
 - Returns NaN if n resolves to NaN or infinity(positive/negative).

Example 12-94 tangent function

```
SELECT tan(n) FROM Archery
```

Value of n	Output
0	0.0
radians(90)	1.633123935319537E16
pi()/4	0.9999999999999999

truncate function

The trunc function truncates the value of an input expression to the specified decimal places on the right of the decimal point.

Syntax:

```
double trunc(n [, d ])
```

Semantics:

- **n, d:** The arguments *n* and *d* are expressions that resolve to numbers. The value of *d* must be in the range of $-30 \leq d \leq 30$. Any value of *d* exceeding the maximum absolute value of 30 is truncated to 30 or -30 respectively. If *d* is not an integer, it is cast to an integer value. For example, `trunc(123.456, 1.5)` is considered `trunc(123.456, 1)`, and `trunc(123.456, 50)` is treated as `trunc(123.456, 30)`.
- **return type:** double.
Depending on the values of *n* and *d*, the `trunc` function behaves as follows:
 - If *d* is a positive value, the function truncates *n* to *d* places on the right of the decimal point, that is, retains only *d* decimal digits to the right of the decimal point.
 - If *d* is greater than or equal the number of digits to the right of the decimal point in *n*, the value of *n* is returned.
 - If *d* is not specified or zero, the function truncates *n* to zero places.
 - If *d* is a negative value, the function truncates *n* to the left of the decimal point, that is, sets *d* digits to the left of the decimal point to zero. The right of the decimal point is also set to zero. If the number of digits in *n* to the left of the decimal point is less than or equal to the absolute value of *d*, then the function returns 0.0.
 - Returns zero if *n* is zero, regardless of the value of *d* (except when *d* is NULL).
 - Returns a NULL value in the following cases:
 - * If either of the arguments is a NULL value
 - * If either of the arguments does not resolve to a number
 - Returns *n* if either of the arguments resolves to a NaN or infinity.

Example 12-95 truncate function

```
SELECT trunc(n,d) FROM Userstocks
```

Value of n	Value of d	Output
111.567	0 or not specified	100.00
111.567	2	100.56
111.567	3	111.567
111.567	-2	100.0
111.567	-3	0.0
-111.567	2	-100.567
-111.567	-1	-110.0

Additional Examples

Learn to use mathematical functions in applications.

Example 12-96 Apply mathematical functions to retrieve intra-day transactions from a trading application

Consider a Userstocks table containing intra-day transactions of a Stock trader in a trading application.

The table DDL is as follows:

```
CREATE TABLE Userstocks (id INTEGER,  
                           stock STRING,  
                           units INTEGER,  
                           buyRate DOUBLE,  
                           sellRate DOUBLE,  
PRIMARY KEY (id))
```

The id field contains the ID of the company, stock field contains the company name in which the user has stock options, the units field contains the number of stocks owned by the user, buyRate is the price at which the user has purchased the stocks, and sellRate is the price at which the user sells his stocks.

Insert sample rows into the table:

```
INSERT into Userstocks VALUES (1, "company1", 100, 10.2, 11.5)  
INSERT into Userstocks VALUES (2, "company2", 20, 15, 14.7)
```

You can use the following query to apply mathematical functions and retrieve the required transaction details from the table:

```
SELECT stock,  
       ceil(units*sellRate-units*buyRate) AS PROFIT,  
       abs(units*sellRate-units*buyRate) AS TURNOVER,  
       round(abs(units*sellRate-units*buyRate)*0.5/100, 2) AS BROKERAGE  
FROM Userstocks
```

Explanation: In the above query, you fetch the day's turnover, profit/loss, and brokerage charges. You calculate the profit/loss using the given arithmetic expressions and apply the [ceil function](#) to get the result as the nearest integer value. You calculate the profit/loss and apply the [abs function](#) to calculate the turnover for each transaction. You calculate the brokerage charges by applying 0.5% on each turnover and round the result to two decimal places using the [round function](#).

Output:

```
{"stock":"company2","PROFIT":-6.0,"TURNOVER":6.0,"BROKERAGE":0.03}  
{"stock":"company1","PROFIT":131.0,"TURNOVER":130.00000000000001,"BROKERAGE":0.65}
```

Example 12-97 Apply mathematical functions to calculate the shooting distance in a gaming application

Consider an Archery table containing data for target shooting in a gaming application.

The table DDL is as follows:

```
CREATE TABLE Archery (sim INTEGER,  
                        angle DOUBLE,  
                        elevation DOUBLE,  
PRIMARY KEY (sim) )
```

The sim field identifies the simulation count, the angle field includes different angles (in degrees) at which a gamer can aim at a target, elevation field contains the height (in meters) at which the target is placed.

Insert sample rows into the table:

```
INSERT INTO Archery VALUES (1, 30, 50)  
INSERT INTO Archery VALUES (2, 45, 50)  
INSERT INTO Archery VALUES (3, 70, 95)
```

You can apply mathematical functions in the following query to calculate the shooting distance for various angles and elevations of the target.

```
SELECT sim, trunc(elevation/sin(radians(angle)),2) AS SLOPE FROM Archery  
ORDER BY  
        sim
```

Explanation: To hit a target placed at a certain elevation, the gamer aims at an angle and shoots a certain distance. These elements form a side of a right-angle triangle, an acute angle, and a hypotenuse. In this query, you calculate the shooting distance for various combinations of target elevation and angles based on the gamer's position from the target. You use the [radians function](#) to convert the angles to radians and calculate the sine value using the [sine function](#). You divide the elevation by the sine value to calculate the distance and truncate the result to two decimal places using the [truncate function](#).

Output:

```
{"sim":1,"SLOPE ":100.0}  
{"sim":2,"SLOPE ":70.71}  
{"sim":3,"SLOPE ":101.09}
```

Example 12-98 Use the mathematical functions to calculate the PH value for a given solution

Consider a PHTable containing Hydrogen ions concentration for different solutions.

The table DDL is as follows:

```
CREATE TABLE PHTable (id INTEGER,  
                        sampleName STRING,  
                        hIons DOUBLE,  
                        phValue DOUBLE,  
PRIMARY KEY (id) )
```

The id field contains the identifier of the solution sample, the sampleName field contains the name of the given sample, the hIons field is the Hydrogen ion concentration in moles per liter of the liquid, and the phValue field is the PH value of the solution. You can use the hIons value

to calculate the pH value of a solution. You initialize the pH value column to NULL while inserting data into the table.

Insert sample rows into the table:

```
INSERT INTO Phtable VALUES (1, "sample1", 0.0063095734448019, NULL)
INSERT INTO Phtable VALUES (2, "sample2", 5.0118723362727E-9, NULL)
INSERT INTO Phtable VALUES (3, "sample3", 1.0E-7, NULL)
```

You can use the mathematical functions to calculate the PH value for a given solution and update the table using the following query:

```
UPDATE Phtable SET pHValue=trunc((log10(1/hIons)),1) where id=1 RETURNING
*
```

Explanation: The PH value indicates the strength of acids or bases in a solution. In this query, you use [log10 function](#) to calculate the PH value of a solution using the formula **PH=log10(1/hIons)**. As the return value of the function is double, you use the [truncate function](#) to truncate the result to one decimal place. You use the [UPDATE Statement](#) to set the resultant PH value in the pH value field for the specified solution.

Output:

```
{"id":1,"sampleName":"sample1","hIons":0.0063095734448019,"pHValue":2.2}
```

You can also use the mathematical functions in the WHERE clause. In this example, you find the acidic samples (PH < 7) using the following query:

```
SELECT sampleName AS ACIDIC FROM Phtable WHERE trunc((log10(1/hIons)),1) <
7.0
```

Output:

```
{"ACIDIC":"sample1"}
```

A

Introduction to the SQL for Oracle NoSQL Database Shell

This appendix describes how to configure, start and use the SQL for Oracle NoSQL Database shell to execute SQL statements. This section also describes the available shell commands.

You can directly execute DDL, DML, user management, security, and informational statements using the SQL shell.

Running the SQL Shell

You can run the SQL shell interactively or use it to run single commands. Here is the general usage to start the shell:

```
java -jar KVHOME/lib/sql.jar
    -helper-hosts <host:port[,host:port]*> -store <storeName>
    [-username <user>] [-security <security-file-path>]
    [-timeout <timeout ms>]
    [-consistency <ABSOLUTE(default) | NONE_REQUIRED>]
    [-durability <COMMIT_SYNC(default) |
                    COMMIT_NO_SYNC | COMMIT_WRITE_NO_SYNC>]
    [single command and arguments]
```

The following are the mandatory parameters:

- `-helper-hosts <host:port[,host:port]*>`: Specifies a comma-separated list of hosts and ports.
- `-store`: Specifies the name of the store.
- `-security`: Specifies the path to the security file in a secure deployment of the store.

For example: `$KVR00T/security/user.security`

The store supports the following optional parameters:

- `-consistency`: Configures the read consistency used for this session. The read operations are serviced either on a master or a replica node depending on the configured value. For more details on consistency, see [Consistency Guarantees](#). The following policies are supported. They are defined in the Consistency class of Java APIs.

If you do not specify this value, the default value `ABSOLUTE` is applied for this session.

- `ABSOLUTE` - The read operation is serviced on a master node. With `ABSOLUTE` consistency, you are guaranteed to obtain the latest updated data.
- `NONE-REQUIRED` - The read operation can be serviced on a replica node. This implies, that if the data is read from the replica node, it may not match what is on the master. However, eventually, it will be consistent with the master.

For more details on the policies, see [Consistency](#) in the *Java Direct Driver API Reference Guide*.

- **-durability**: Configures the write durability setting used in this session. This value defines the durability policies to be applied for achieving master commit synchronization, that is, the actions performed by the master node to return with a normal status from the write operations. For more details on durability, see [Durability Guarantees](#).

If you do not specify this value, the default value `COMMIT_SYNC` is applied for this session.

- `COMMIT_NO_SYNC` - The data is written to the host's in-memory cache, but the master node does not wait for the data to be written to the file system's data buffers or subsequent physical storage.
- `COMMIT_SYNC` - The data is written to the in-memory cache, transferred to the file system's data buffers, and then synchronized to a stable storage before the write operation completes normally.
- `COMMIT_WRITE_NO_SYNC` - The data is written to the in-memory cache, and transferred to the file system's data buffers, but not necessarily into physical storage.

For more details on the policies, see [Durability](#) in the *Java Direct Driver API Reference Guide*.

- **-timeout**: Configures the request timeout used for this session. The default value is 5000ms.
- **-username**: Specifies a username to log on as in a secure deployment.

For example, you can start the shell as follows:

```
java -jar KVHOME/lib/sql.jar  
-helper-hosts node01:5000 -store kvstore  
sql->
```

This command assumes that a store `kvstore` is running at port 5000. After the SQL starts successfully, you execute queries. In the next part of this document, you will find an introduction to SQL for Oracle NoSQL Database and how to create query statements.

If you want to import records from a file in either JSON or CSV format, you can use the `import` command. For more information see [import](#).

If you want to run a script, use the `load` command. For more information see [load](#).

```
sql-> command [arguments]
```

-single command and arguments: Specifies the utility commands that can be accessed from the SQL shell. You can use them with the syntax shown above.

For a complete list of utility commands accessed through "java -jar" `<kvhome>/lib/sql.jar <command>` see [Shell Utility Commands](#).

Configuring the shell

You can also set the shell start-up arguments by modifying the user-specific configuration file `.kvclirc`, typically found in your home directory (for example, `~/.kvclirc`). If the file doesn't already exist, you may need to create it manually.

You configure the arguments in the `.kvclirc` file using the `name=value` format. This file is shared by all shells, each having its named section. `[sql]` is used for the Query shell, while `[kvcli]` is used for the Admin Command Line Interface (CLI).

For example, the `.kvclirc` file would then contain content as follows:

```
[sql]
helper-hosts=node01:5000
store=kvstore
timeout=10000
consistency=NONE_REQUIRED
durability=COMMIT_NO_SYNC
username=root
security=/tmp/login_root

[kvcli]
host=node01
port=5000
store=kvstore
admin-host=node01
admin-port=5001
username=user1
security=/tmp/login_user
admin-username=root
admin-security=/tmp/login_root
timeout=10000
consistency=NONE_REQUIRED
durability=COMMIT_NO_SYNC
```

Shell Utility Commands

The following sections describe the utility commands accessed through "java -jar" `<kvhome>/lib/sql.jar <command>`.

The interactive prompt for the shell is:

```
sql->
```

The shell comprises a number of commands. All commands accept the following flags:

- `-help`
Displays online help for the command.
- `?`
Synonymous with `-help`. Displays online help for the command.

The shell commands have the following general format:

1. All commands are structured like this:

```
sql-> command [arguments]
```

2. All arguments are specified using flags that start with `"-"`
3. Commands and subcommands are case-insensitive and match on partial strings(prefixes) if possible. The arguments, however, are case-sensitive.
4. All commands must terminate with a semicolon `;"`

connect

```
connect -host <hostname> -port <port> -name <storeName>
        [-timeout <timeout ms>]
        [-consistency <ABSOLUTE(default) | NONE_REQUIRED>]
        [-durability <COMMIT_SYNC(default) | COMMIT_NO_SYNC |
COMMIT_WRITE_NO_SYNC>]
        [-username <user>] [-security <security-file-path>]
```

Connects to a data store to perform data access functions. If the instance is secured, you may need to provide the log in credentials.

The following are the mandatory parameters:

- -host: Identifies the host name of a node in your store.
- -port <port>: The TCP/IP port on which Oracle NoSQL Database should be contacted. This port must be free (unused) on each node.
- -name <storename>: Specifies the name of the store.
- -security: Specifies the path to the security file in a secure deployment of the store.

For example: \$KVR00T/security/user.security

The store supports the following optional parameters:

- -consistency: Configures the read consistency used for this session. The read operations are serviced either on a master or a replica node depending on the configured value. For more details on consistency, see [Consistency Guarantees](#). The following policies are supported. They are defined in the Consistency class of Java APIs.

If you do not specify this value, the default value ABSOLUTE is applied for this session.

- ABSOLUTE - The read operation is serviced on a master node. With ABSOLUTE consistency, you are guaranteed to obtain the latest updated data.
- NONE-REQUIRED - The read operation can be serviced on a replica node. This implies, that if the data is read from the replica node, it may not match what is on the master. However, eventually, it will be consistent with the master.

For more details on the policies, see [Consistency](#) in the *Java Direct Driver API Reference Guide*.

- -durability: Configures the write durability setting used in this session. This value defines the durability policies to be applied for achieving master commit synchronization, that is, the actions performed by the master node to return with a normal status from the write operations. For more details on durability, see [Durability Guarantees](#).

If you do not specify this value, the default value COMMIT_SYNC is applied for this session.

- COMMIT_NO_SYNC - The data is written to the host's in-memory cache, but the master node does not wait for the data to be written to the file system's data buffers or subsequent physical storage.
- COMMIT_SYNC - The data is written to the in-memory cache, transferred to the file system's data buffers, and then synchronized to a stable storage before the write operation completes normally.
- COMMIT_WRITE_NO_SYNC - The data is written to the in-memory cache, and transferred to the file system's data buffers, but not necessarily into physical storage.

For more details on the policies, see [Durability](#) in the *Java Direct Driver API Reference Guide*.

- -timeout: Specifies the store request timeout in milliseconds.
- -username: Specifies a username to log on as in a secure deployment.

consistency

```
consistency [[ABSOLUTE(default) | NONE_REQUIRED] [-time -permissible-lag
<time_ms> -timeout <time_ms>]]
```

Configures the read consistency used for this session.

The read operations are serviced either on a master or a replica node depending on the configured value. For more details on consistency, see Consistency Guarantees. The following policies are supported. They are defined in the Consistency class of Java APIs.

If you do not specify this value, the default value ABSOLUTE is applied for this session.

- ABSOLUTE - The read operation is serviced on a master node. With ABSOLUTE consistency, you are guaranteed to obtain the latest updated data.
- NONE-REQUIRED - The read operation can be serviced on a replica node. This implies, that if the data is read from the replica node, it may not match what is on the master. However, eventually, it will be consistent with the master.

For more details on the policies, see [Consistency](#) in the *Java Direct Driver API Reference Guide*.

Other non-mandatory parameter includes:

-time: Indicates the use of time-based options for tuning the consistency. The following parameters are supported:

- -permissible-lag: Sets the maximum allowable delay between the replica node and master. Enter the value in milliseconds.
- -timeout: Configures the request timeout used for this session. The default value is 5000ms.

For example, you can define the consistency as follows:

```
sql-> consistency ABSOLUTE -time -permissible-lag 1000 -timeout 5000;
```

The following sample output confirms the applied read policy:

```
Read consistency policy: Consistency.Time[permissibleLag_ms=1000,
timeout_ms=5000]
```

describe

```
describe | desc [as json]
{table table_name [field_name[,...]] |
index index_name on table_name
}
```

Describes information about a table or index, optionally in JSON format.

Specify a fully-qualified *table_name* as follows:

Entry specification	Description
<i>table_name</i>	Required. Specifies the full table name. Without further qualification, this entry indicates a table created in the default namespace (sysdefault), which you do not have to specify.
<i>parent-table.child-table</i>	Specifies a child table of a parent. Specify the parent table followed by a period (.) before the child name. For example, if the parent table is Users, specify the child table named MailingAddress as Users.MailingAddress.
<i>namespace-name:table-name</i>	Specifies a table created in the non-default namespace. Use the namespace followed by a colon (:). For example, to reference table Users, created in the Sales namespace, enter <i>table_name</i> as Sales:Users.

Following is the output of describe for table ns1:t1:

```
sql-> describe table ns1:t1;
=== Information ===
+-----+-----+-----+-----+-----+-----+-----+-----+
| namespace | name | ttl | owner | sysTable | r2compat | parent | children |
indexes | description |
+-----+-----+-----+-----+-----+-----+-----+-----+
| ns1      | t1   |    |      | N        | N        |      |      |
|          |      |    |      |          |          |      |      |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+

=== Fields ===
+---+-----+-----+-----+-----+-----+-----+-----+
+-----+
| id | name | type  | nullable | default | shardKey | primaryKey |
identity |
+---+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | id   | Integer | N        | NullValue | Y        | Y        |
|   |      |         |          |           |          |          |
+---+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2 | name | String  | Y        | NullValue |          |          |
|   |      |         |          |           |          |          |
+---+-----+-----+-----+-----+-----+-----+-----+
+-----+

sql->
```

This example shows using `describe as json` for the same table:

```
sql-> describe as json table ns1:t1;
{
  "json_version" : 2,
  "type" : "table",
  "namespace" : "ns1",
  "name" : "t1",
  "id" : "55",
  "fields" : [{
    "name" : "id",
    "type" : "INTEGER",
    "nullable" : false
  }, {
    "name" : "name",
    "type" : "STRING",
    "nullable" : true
  }],
  "primaryKey" : ["id"],
  "shardKey" : ["id"]
}
```

durability

```
durability [[COMMIT_WRITE_NO_SYNC | COMMIT_SYNC |
COMMIT_NO_SYNC] | [-master-sync <sync-policy> -replica-sync <sync-policy>
-replica-ask <ack-policy>]] <sync-policy>: SYNC, NO_SYNC, WRITE_NO_SYNC
<ack-policy>: ALL, NONE, SIMPLE_MAJORITY
```

Configures the write durability used for this session.

exit

```
exit | quit
```

Exits the interactive command shell.

help

```
help [command]
```

Displays help message for all shell commands and `sql` command.

history

```
history [-last <n>] [-from <n>] [-to <n>]
```

Displays command history. By default all history is displayed. Optional flags are used to choose ranges for display.

import

```
import -table table_name -file file_name [JSON | CSV]
```

Imports records from the specified file into table *table_name*.

Specify a fully-qualified *table_name* as follows:

Entry specification	Description
<i>table_name</i>	Required. Specifies the full table name. Without further qualification, this entry indicates a table created in the default namespace (sysdefault), which you do not have to specify.
<i>parent-table.child-table</i>	Specifies a child table of a parent. Specify the parent table followed by a period (.) before the child name. For example, if the parent table is Users, specify the child table named MailingAddress as Users.MailingAddress.
<i>namespace-name:table-name</i>	Specifies a table created in the non-default namespace. Use the namespace followed by a colon (:). For example, to reference table Users, created in the Sales namespace, enter <i>table_name</i> as Sales:Users.

Use -table to specify the name of a table into which the records are loaded. The alternative way to specify the table is to add the table specification "Table: *table_name*" before its records in the file.

For example, this file contains the records to insert into two tables, users and email:

```
Table: users
<records of users>
...
Table: emails
<record of emails>
...
```

The imported records can be either in JSON or CSV format. If you do not specify the format, JSON is assumed.

load

```
load -file <path to file>
```

Load the named file and interpret its contents as a script of commands to be executed. If any command in the script fails execution will end.

For example, suppose the following commands are collected in the script file test.sql:

```
### Begin Script ###
load -file test.ddl
```

```
import -table users -file users.json
### End Script ###
```

Where the file `test.ddl` would contain content like this:

```
DROP TABLE IF EXISTS users;
CREATE TABLE users(id INTEGER, firstname STRING, lastname STRING,
age INTEGER, primary key (id));
```

And the file `users.json` would contain content like this:

```
{"id":1,"firstname":"Dean","lastname":"Morrison","age":51}
{"id":2,"firstname":"Idona","lastname":"Roman","age":36}
{"id":3,"firstname":"Bruno","lastname":"Nunez","age":49}
```

Then, the script can be run by using the `load` command in the shell:

```
> java -jar KVHOME/lib/sql.jar -helper-hosts node01:5000 \
-store kvstore
sql-> load -file ./test.sql
Statement completed successfully.
Statement completed successfully.
Loaded 3 rows to users.
```

mode

```
mode [COLUMN | LINE | JSON [-pretty] | CSV]
```

Sets the output mode of query results. The default value is JSON.

For example, a table shown in COLUMN mode:

```
sql-> mode column;
sql-> SELECT * from users;
```

id	firstname	lastname	age
8	Len	Aguirre	42
10	Montana	Maldonado	40
24	Chandler	Oneal	25
30	Pascale	Mcdonald	35
34	Xanthus	Jensen	55
35	Ursula	Dudley	32
39	Alan	Chang	40
6	Lionel	Church	30
25	Alyssa	Guerrero	43
33	Gannon	Bray	24
48	Ramona	Bass	43
76	Maxwell	Mcleod	26
82	Regina	Tillman	58
96	Iola	Herring	31
100	Keane	Sherman	23

```
+-----+-----+-----+-----+
...
```

100 rows returned

Empty strings are displayed as an empty cell.

```
sql-> mode column;
sql-> SELECT * from tab1 where id = 1;
+----+-----+-----+-----+
| id | s1 | s2 | s3 |
+----+-----+-----+-----+
| 1 | NULL |  | NULL |
+----+-----+-----+-----+
```

1 row returned

For nested tables, indentation is used to indicate the nesting under column mode:

```
sql-> SELECT * from nested;
+---+-----+-----+-----+-----+
| id | name | details |
+---+-----+-----+-----+
| 1 | one | address |
|   |   | city    | Waitakere |
|   |   | country | French Guiana |
|   |   | zipcode | 7229 |
|   |   | attributes |
|   |   | color    | blue |
|   |   | price    | expensive |
|   |   | size     | large |
|   |   | phone    | [(08)2435-0742, (09)8083-8862, (08)0742-2526] |
+---+-----+-----+-----+-----+
| 3 | three | address |
|   |   | city    | Viddalba |
|   |   | country | Bhutan |
|   |   | zipcode | 280071 |
|   |   | attributes |
|   |   | color    | blue |
|   |   | price    | cheap |
|   |   | size     | small |
|   |   | phone    | [(08)5361-2051, (03)5502-9721, (09)7962-8693] |
+---+-----+-----+-----+-----+
...
```

For example, a table shown in LINE mode, where the result is displayed vertically and one value is shown per line:

```
sql-> mode line;
sql-> SELECT * from users;

> Row 1
+-----+-----+
| id | 8 |
+-----+-----+
```

```

| firstname | Len      |
| lastname  | Aguirre  |
| age       | 42       |
+-----+-----+

> Row 2
+-----+-----+
| id        | 10       |
| firstname | Montana  |
| lastname  | Maldonado|
| age       | 40       |
+-----+-----+

> Row 3
+-----+-----+
| id        | 24       |
| firstname | Chandler  |
| lastname  | Oneal    |
| age       | 25       |
+-----+-----+

...
100 rows returned

```

Just as in COLUMN mode, empty strings are displayed as an empty cell:

```

sql-> mode line;
sql-> SELECT * from tab1 where id = 1;

> Row 1
+-----+-----+
| id      | 1      |
| s1      | NULL   |
| s2      |        |
| s3      | NULL   |
+-----+-----+

1 row returned

```

For example, a table shown in JSON mode:

```

sql-> mode json;
sql-> SELECT * from users;
{"id":8,"firstname":"Len","lastname":"Aguirre","age":42}
{"id":10,"firstname":"Montana","lastname":"Maldonado","age":40}
{"id":24,"firstname":"Chandler","lastname":"Oneal","age":25}
{"id":30,"firstname":"Pascale","lastname":"Mcdonald","age":35}
{"id":34,"firstname":"Xanthus","lastname":"Jensen","age":55}
{"id":35,"firstname":"Ursula","lastname":"Dudley","age":32}
{"id":39,"firstname":"Alan","lastname":"Chang","age":40}
{"id":6,"firstname":"Lionel","lastname":"Church","age":30}
{"id":25,"firstname":"Alyssa","lastname":"Guerrero","age":43}
{"id":33,"firstname":"Gannon","lastname":"Bray","age":24}
{"id":48,"firstname":"Ramona","lastname":"Bass","age":43}
{"id":76,"firstname":"Maxwell","lastname":"Mcleod","age":26}
{"id":82,"firstname":"Regina","lastname":"Tillman","age":58}

```



```
{"id":96,"firstname":"Iola","lastname":"Herring","age":31}
{"id":100,"firstname":"Keane","lastname":"Sherman","age":23}
{"id":3,"firstname":"Bruno","lastname":"Nunez","age":49}
{"id":14,"firstname":"Thomas","lastname":"Wallace","age":48}
{"id":41,"firstname":"Vivien","lastname":"Hahn","age":47}
...
100 rows returned
```

Empty strings are displayed as "".

```
sql-> mode json;
sql-> SELECT * from tab1 where id = 1;
{"id":1,"s1":null,"s2":"","s3":"NULL"}
```

1 row returned

Finally, a table shown in CSV mode:

```
sql-> mode csv;
sql-> SELECT * from users;
8,Len,Aguirre,42
10,Montana,Maldonado,40
24,Chandler,Oneal,25
30,Pascale,Mcdonald,35
34,Xanthus,Jensen,55
35,Ursula,Dudley,32
39,Alan,Chang,40
6,Lionel,Church,30
25,Alyssa,Guerrero,43
33,Gannon,Bray,24
48,Ramona,Bass,43
76,Maxwell,McLeod,26
82,Regina,Tillman,58
96,Iola,Herring,31
100,Keane,Sherman,23
3,Bruno,Nunez,49
14,Thomas,Wallace,48
41,Vivien,Hahn,47
...
100 rows returned
```

Like in JSON mode, empty strings are displayed as "".

```
sql-> mode csv;
sql-> SELECT * from tab1 where id = 1;
1,NULL,"","NULL"
```

1 row returned

Note

Only rows that contain simple type values can be displayed in CSV format. Nested values are not supported.

output

`output [stdout | file]`

Enables or disables output of query results to a file. If no argument is specified, it shows the current output.

page

`page [on | <n> | off]`

Turns query output paging on or off. If specified, *n* is used as the page height.

If *n* is 0, or "on" is specified, the default page height is used. Setting *n* to "off" turns paging off.

show faults

`show faults [-last] [-command <index>]`

Encapsulates commands that display the state of the store and its components.

show ddl

`show ddl <table>`

The `show ddl` query retrieves the DDL statement for a specified table. If the table has indexes, the statement returns the DDLs for the table and the indexes.

Example : Fetch the DDL for a specified table.

The following statement fetches the DDL for the `BaggageInfo` table.

```
show ddl BaggageInfo;
```

Output:

```
CREATE TABLE IF NOT EXISTS BaggageInfo (ticketNo LONG, fullName STRING,  
gender STRING,  
contactPhone STRING, confNo STRING, bagInfo JSON, PRIMARY  
KEY(SHARD(ticketNo)))
```

In the following example, the `fixedschema_contact` index exists in the `BaggageInfo` table. The statement retrieves the DDLs for the `BaggageInfo` table and `fixedschema_contact` index on the table.

```
show ddl BaggageInfo;
```

Output:

```
CREATE TABLE IF NOT EXISTS BaggageInfo (ticketNo LONG, fullName STRING,  
gender STRING,  
contactPhone STRING, confNo STRING, bagInfo JSON, PRIMARY  
KEY(SHARD(ticketNo)))CREATE INDEX IF NOT EXISTS fixedschema_contact ON  
BaggageInfo(contactPhone)
```

show indexes

```
show_indexes_statement ::= SHOW [AS JSON] INDEXES ON table_name
```

The `show indexes` statement provides the list of indexes present on a specified table. The parameter `AS JSON` is optional and can be specified if you want the output to be in JSON format.

Example 1: List indexes on the specified table

The following statement lists the indexes present on the `users2` table.

```
SHOW INDEXES ON users2;  
indexes  
  idx1
```

Example 2: List indexes on the specified table in JSON format

The following statement lists the indexes present on the `users2` table in JSON format.

```
SHOW AS JSON INDEXES ON users2;  
{  
  "indexes" :  
    ["idx1"]  
}
```

show namespaces

```
show [AS JSON] namespaces
```

Shows a list of all namespaces in the system.

For example:

```
sql-> show namespaces  
namespaces  
  ns1
```

```
sysdefault
sql-> show as json namespaces
{"namespaces" : ["ns1","sysdefault"]}
```

show query

```
show query <statement>
```

Displays the query plan for a query.

For example:

```
sql-> show query SELECT * from Users;
RECV([6], 0, 1, 2, 3, 4)
[
  DistributionKind : ALL_PARTITIONS,
  Number of Registers :7,
  Number of Iterators :12,
  SFW([6], 0, 1, 2, 3, 4)
  [
    FROM:
    BASE_TABLE([5], 0, 1, 2, 3, 4)
    [Users via primary index] as $$Users

    SELECT:
    *
  ]
]
```

show regions

```
show_regions_statement ::= SHOW [AS JSON] REGIONS
```

The `show regions` statement provides the list of regions present in a multi-region Oracle NoSQL Database setup. The parameter `AS JSON` is optional and can be specified if you want the output to be in JSON format.

Example 1: Fetching all regions in a multi-region database setup

```
SHOW REGIONS;
regions
  my_region1 (remote, active)
  my_region2 (remote, active)
```

Example 2: Fetching all regions in a multi-region database setup in JSON format

```
SHOW AS JSON REGIONS;
{"regions" : [
  {"name" : "my_region1", "type" : "remote", "state" : "active"},
  {"name" : "my_region2", "type" : "remote", "state" : "active"}
]}
```

show roles

```
show [as json] roles | role <role_name>
```

Shows either all the roles currently defined for the store, or the named role.

show tables

```
show [as json] {tables | table table_name}
```

Shows either all tables in the data store, or one specific table, *table_name*.

Specify a fully-qualified *table_name* as follows:

Entry specification	Description
<i>table_name</i>	Required. Specifies the full table name. Without further qualification, this entry indicates a table created in the default namespace (sysdefault), which you do not have to specify.
<i>parent-table.child-table</i>	Specifies a child table of a parent. Specify the parent table followed by a period (.) before the child name. For example, if the parent table is Users, specify the child table named MailingAddress as Users.MailingAddress.
<i>namespace-name:table-name</i>	Specifies a table created in the non-default namespace. Use the namespace followed by a colon (:). For example, to reference table Users, created in the Sales namespace, enter <i>table_name</i> as Sales:Users.

The following example indicates how to list all tables, or just one table. The empty tableHierarchy field indicates that table t1 was created in the default namespace:

```
sql-> show tables
tables
  SYS$IndexStatsLease
  SYS$PartitionStatsLease
  SYS$SGAttributesTable
  SYS$TableStatsIndex
  SYS$TableStatsPartition
  ns10:t10
  parent
  parent.child
  sg1
  t1

sql-> show table t1
tableHierarchy
  t1
```

To show a table created in a namespace, as shown in the list of all tables, fully-qualify *table_name* as follows. In this case, tableHierarchy field lists namespace ns1 in which table t1 was created. The example also shows how the table is presented as json:

```
sql-> show tables;
tables
  SYS$IndexStatsLease
  SYS$PartitionStatsLease
  SYS$SGAttributesTable
  SYS$TableStatsIndex
  SYS$TableStatsPartition
  ns1:foo
  ns1:t1

sql-> show table ns1:t1;
tableHierarchy(namespace ns1)
  t1
sql-> show as json table ns1:t1;
{"namespace": "ns1"
"tableHierarchy" : ["t1"]}
```

show users

```
show [as json] users | user <user_name>
```

Shows either all the users currently existing in the store, or the named user.

timeout

```
timeout [<timeout_ms>]
```

The timeout command configures or displays the request timeout for this session in milliseconds(ms).

The request timeout is the amount of time that the client will wait to get a response to a request that it has sent.

If the optional timeout_ms attribute is specified, then the request timeout is set to the specified value.

If the optional timeout_ms attribute is not specified, then the current value of request timeout is displayed.

Example A-1 timeout

The following example gets the current value of the request timeout.

```
sql-> timeout
Request timeout used: 5,000ms
```

Example A-2 timeout

The following example set the request timeout value to 20000 milliseconds (20 seconds).

```
sql-> timeout 20000
Request timeout used: 20,000ms
```

Note

A shell command may require multiple requests to a server or servers. The timeout applies to each such individual request. A shell command sends out multiple requests and has to wait for each of them to return before the command is finished. As a result, a shell command may have to wait for longer time than the specified timeout and this total wait could be greater than the wait time of the individual request.

timer

```
timer [on | off]
```

Turns the measurement and display of execution time for commands on or off. If not specified, it shows the current state of timer. For example:

```
sql-> timer on
sql-> SELECT * from users where id <= 10 ;
+---+-----+-----+-----+
| id | firstname | lastname | age |
+---+-----+-----+-----+
| 8 | Len      | Aguirre  | 42 |
| 10 | Montana  | Maldonado | 40 |
| 6 | Lionel   | Church   | 30 |
| 3 | Bruno    | Nunez    | 49 |
| 2 | Idona    | Roman    | 36 |
| 4 | Cooper   | Morgan   | 39 |
| 7 | Hanae    | Chapman  | 50 |
| 9 | Julie    | Taylor   | 38 |
| 1 | Dean     | Morrison | 51 |
| 5 | Troy     | Stuart   | 30 |
+---+-----+-----+-----+
```

10 rows returned

Time: 0sec 98ms

verbose

```
verbose [on | off]
```

Toggles or sets the global verbosity setting. This property can also be set on a per-command basis using the -verbose flag.

version

version

Display client version information.