Oracle® NoSQL Database

Developers Guide

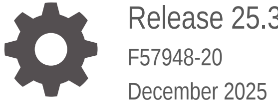

Release 25.3
F57948-20
December 2025

Oracle NoSQL Database Developers Guide, Release 25.3

F57948-20

Copyright © 2022, 2025, 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 Get Started

Getting started with Oracle NoSQL Database 1
- Start your data store 2
- Starting the SQL shell 2
- Sample use-cases used in the examples 3
- Tables used in the examples 4
- Describe tables 5
- Sample data to run queries 7
- Table Hierarchies 13
- About Oracle NoSQL Database SDK drivers 14
- Obtaining a NoSQL Handle 17

2 Create

Creating a namespace 1

Creating a region 3

Creating a table 8
- Using SQL commands 9
- Using TableRequest API 14

Create and View Indexes 20
- Classification of Indexes 21
- Creating Indexes 22
  - Using SQL commands 23
  - Using TableRequest API 29
- View Index 31

3 Manage

Namespace Management 1
- Namespace Resolution 1
- Manage Namespaces 1
- Namespace scoped privileges 4
- Granting Authorization Access to Namespaces 4

Managing Tables, Indexes & Regions 7

Copyright © 2022, 2025, Oracle and/or its affiliates.

Alter Table 7
Using SQL command to alter table 8
Using TableRequest API to alter table 10
Drop Table 12
Using SQL command to drop table 12
Using TableRequest API to drop table 13
Drop Index 15
Using SQL command to drop index 15
Using TableRequest API to drop index 15
Manage regions 17

4 Develop

Inserting, Modifying, and Deleting Data 1
Insert data 1
Using SQL command to insert data 1
Using Put API to insert data 7
Using MultiWrite API to insert data 12
Upsert Data 15
Using SQL command to upsert data 15
Using API to upsert data 23
Update Data 36
Using SQL command to update data 36
Using API to update data 44
Modify JSON data 47
Using SQL command 47
Using API 48
Delete Data 52
Using SQL command to delete data 53
Using API to delete a single row 53
Using API to delete multiple rows 57
Using Query API to delete data 60
Simple SELECT queries 63
Using SQL commands to fetch data 64
Substituting column names in a query 70
Using Get API to fetch data 72
Using Query API to fetch data 76
SELECT queries on JSON collection tables 82
Using Path expressions 84
Using Internal variables and aliases 84
Working with Arrays 85
Working with nested data type 87

Copyright © 2022, 2025, Oracle and/or its affiliates.

Finding the size of a complex data type 88
Using user-defined row metadata 89
Using row metadata in Write Operations 90
Using row metadata in Read Operations 91
Using SQL commands on row metadata 92
Using Left Outer joins with parent-child tables 93
Overview of Left Outer Joins 93
Examples using Left Outer Joins 94
SQL Examples 95
Query API examples 102
Using NESTED TABLES to join parent-child tables 105
Overview of NESTED TABLES 105
Examples using NESTED TABLES 106
SQL Examples 107
Query API examples 114
Using inner join with parent-child tables 118
Overview of Inner Join 118
Examples using Inner Join 119
SQL Examples 121
Query API Examples 124
Tuning and Optimizing SQL queries 127
Using Indexes for query optimization 128
Examples of queries using index 128
Managing GeoJSON data 135
geo_inside 136
geo_intersect 137
geo_distance 139
geo_within_distance 140
geo_near 141
geo_is_geometry 143

5 Reference

Operators in SQL 1
Sequence Comparison Operators 1
Logical operators 3
NULL operators 5
Value Comparison Operators 6
BETWEEN Operator 9
IN Operator 10
Regular Expression Conditions 11
EXISTS Operator 12

Copyright © 2022, 2025, Oracle and/or its affiliates.

Is-Of-Type Operator 13
SQL Operators examples using QueryRequest API 14
Sorting, Grouping & Limiting results 20
Ordering results 20
Limit and offset results 21
Grouping results 23
Aggregating results 24
Examples using QueryRequest API 25
Primary Expressions in SQL 31
Parenthesized Expressions 31
Case Expressions 31
Cast Expression 33
Sequence Transform Expressions 35
Extract Expressions 36
SQL Expression examples using QueryRequest API 38
Timestamp Functions 45
Timestamp Arithmetic Functions 46
Timestamp Round Functions 51
Timestamp Format Functions 53
Timestamp Extract Functions 54
Current Time Functions 56
Examples using QueryRequest API 57
Functions on Strings 61
substring function 62
concat function 62
upper and lower functions 63
trim function 63
length function 64
contains function 65
starts_with and ends_with functions 65
index_of function 66
replace function 67
reverse function 68
Examples using QueryRequest API 68
Query execution plan 73
Overview of query plan 73
Query 1: Using primary key index with an index range scan 76
Query 2: Using primary key index with an index predicate 78
Query 3: Using a secondary index with an index range scan 82
Query 4: Using the primary index 84
Query 5: Sort the data using a Covering index 86
Query 6: Using a secondary index with an index predicate 88

Copyright © 2022, 2025, Oracle and/or its affiliates.

Query 7: Group data with fields as part of the index 91
Query 8: Using the secondary index with multiple index scans 93
Query 9: A SINGLE PARTITION query using a primary index 96
Query 10: Group data with fields not part of any index 99
Table Modelling and Design 102
Schema Flexibility in Oracle NoSQL Database 103
Choice of Keys in NoSQL Database 105
Using Indexes in NoSQL Database 108
Transactions in NoSQL database 109
Handling Errors 111
Handling Driver Errors 112

Index

Copyright © 2022, 2025, Oracle and/or its affiliates.

List of Tables

3-1	Namespace Privileges and Permissions	4
4-1	Nested Tables Vs LOJ	106
5-1	Timestamp functions	45
5-2	Comparison between Identity Column and UUID column	108

Copyright © 2022, 2025, Oracle and/or its affiliates.

1

Get Started

The articles in this section focus on providing the quickest path to use Oracle NoSQL Database. It contains the details to connect to the database, details of the schemas used in the examples, and sample data to run queries.

Getting started with Oracle NoSQL Database

Oracle NoSQL Database is a distributed, shared-nothing, non-relational database that provides large-scale storage and access to key/value, JSON, and tabular data. It can deliver predictable, low latencies to simple queries at any scale and is designed from the ground up for high availability.

Oracle NoSQL Database offers highly flexible deployment options. In most scenarios, it is deployed on a cluster of commodity computers connected by a high-speed network. Oracle NoSQL Database offers various methods by which your application can access the database:

- You can use the SDKs available in different programming languages to develop your applications. Oracle NoSQL Database offers two types of SDKs:
  - **Direct Driver SDK:** This SDK enables applications to directly connect with the Oracle NoSQL Database node using TCP/IP. Hence, you must ensure that a network route exists between the application and every Oracle NoSQL node in the database cluster. The only supported programming language for the direct driver is Java .
  - **Standard SDKs:** These SDKs enable applications to connect to the database using HTTP protocol via the Oracle NoSQL HTTP proxy.
    Oracle NoSQL Database supports many of the most popular programming languages and frameworks with idiomatic language APIs and data structures, giving your application access to data stored in the database. It currently supports the following: Java, Python, Node.js (JavaScript/TypeScript), Golang, C#/.NET, Spring and Rust.

    You can refer to the table below for links to the SDKs, API guides, and examples:

    * SDK (GitHub) - Provides details on how to install, connect and get started with the SDK
    * API Guide - Provides the packages, classes, methods, and interfaces available in the SDK
    * Examples - Provides code samples that you can try out

SDK (GitHub)	API Guide	Examples
Oracle NoSQL Java SDK	Java SDK API Guide	Java Examples
Oracle NoSQL Python SDK	Python SDK API Guide	Python Examples
Oracle NoSQL Go SDK	Go SDK API Guide	Go Examples
Oracle NoSQL Node.js SDK	Node.js SDK API Guide	Node.js Examples
Oracle NoSQL .NET SDK	.NET SDK API Reference	.NET Examples
Oracle NoSQL Rust SDK	Rust SDK API Guide	Rust Examples
Oracle NoSQL Spring SDK	Spring SDK API Guide	Spring Examples

Copyright © 2022, 2025, Oracle and/or its affiliates.

- You can use integrated development environments to develop your software efficiently. Oracle NoSQL Database plugins are available for the Visual Studio Code and IntelliJ IDEs.
- You can use the SQL language in your applications to interact with the Oracle NoSQL Database.

Start your data store

You can quickly get started using KVLite, which is a simplified version of the Oracle NoSQL Database

KVLite provides a single storage node, single shard store, that is not replicated. It runs in a single process without requiring any administrative interface. The Oracle NoSQL Database Proxy is a middle-tier component that lets any application written using one of the Oracle NoSQL Database SDKs communicate with the Oracle NoSQL Database. The JAR file for the Oracle NoSQL Database Proxy is included in the Enterprise Edition distribution and the Community Edition distribution of Oracle NoSQL Database. You can quickly start your data store by following the two steps below.

1. Install KVLite
2. Start KVLite

Starting the SQL shell

You can run SQL queries and DDL statements directly from the SQL shell. 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`: 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: `$KVROOT/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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- 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 the username to log in as.

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 run queries.

```
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 details on supported shell utility commands, see Shell Utility Commands in the *SQL Beginner's Guide*.

Sample use-cases used in the examples

You have two different schemas (with real-time scenarios) for learning various SQL concepts. These two schemas will include various data types that can be used in the Oracle NoSQL database.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Schema 1: BaggageInfo schema

Using this schema you can handle a use case wherein passengers traveling on a flight can track the progress of their checked-in bags or luggage along the route to the final destination. This functionality can be made available as part of the airline's mobile application. Once the passenger logs into the mobile application, the ticket number or reservation code of the current flight is displayed on the screen. Passengers can use this information to search for their baggage information. The mobile application is using NoSQL Database to store all the data related to the baggage. In the backend, the mobile application logic performs SQL queries to retrieve the required data.

Schema 2: Streaming Media Service - Persistent User Profile Store

Consider a TV streaming application. It streams various shows that are watched by customers across the globe. Every show has a number of seasons and every season has multiple episodes. You need a persistent meta-data store that keeps track of the current activity of the customers using the TV streaming application. Using this schema you can provide useful information to the customer such as episodes they watched, the watch time per episode, the total number of seasons of the show they watched, etc. The data is stored in the NoSQL Database and the application performs SQL queries to retrieve the required data and make it available to the user.

Tables used in the examples

The table is the basic structure to hold user data.

Table 1: Airline baggage tracking application

The table used in this schema is `BaggageInfo`. This schema has a combination of fixed data types like LONG, STRING. It also has a schema-less JSON (`bagInfo`) as one of its columns. The schema-less JSON does not have a fixed data type. The bag information of the passengers is a schema-less JSON. In contrast, the passenger's information like ticket number, full name, gender, contact details is all part of a fixed schema. You can add any number of fields to this non-fixed schemaless JSON field. .

The following code creates the table.

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

Table 2: Streaming Media Service - Persistent user profile store

The table used in this schema is `stream_acct`. The primary key in this schema is `acct_id`. The schema also includes a JSON column (`acct_data`), which is schema-less. The schema-less JSON does not have a fixed data type. You can add any number of fields to this non-fixed schema-less JSON field.

Copyright © 2022, 2025, Oracle and/or its affiliates.

The following code creates the table.

```
CREATE TABLE stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data JSON,
PRIMARY KEY(acct_id)
)
```

Table 3: JSON collection table - Shopping application

JSON collection tables are useful for applications that store and retrieve data purely as documents. JSON collection tables are schema-less tables, which provide the flexibility to create tables with primary key field declaration. You must supply the value of primary key fields along with the other fields in the document during the insertion of data into the table.

The table used in shopping application is `storeAcct`. 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 creates the table:

```
CREATE TABLE IF NOT EXISTS storeAcct (
contactPhone STRING,
PRIMARY KEY(SHARD(contactPhone))
) AS JSON COLLECTION
```

Describe tables

You use DESCRIBE or DESC command to view the description of a table.

```
(DESCRIBE | DESC) [AS JSON] TABLE table_name [ "(" field_name")"]
```

AS JSON can be specified if you want the output to be in JSON format. You could get information about a specific field in any table by providing the field name.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 1: Describe a table

```
DESCRIBE TABLE stream_acct
```

Output:

```
 === Information ===
 +-------------+-----+-------+----------------+----------+--------+----------
+---------+---------+-------------+
 |    name     | ttl | owner | jsonCollection | sysTable | parent | children
| regions | indexes | description |
 +-------------+-----+-------+----------------+----------+--------+----------
+---------+---------+-------------+
 | stream_acct |     |       | N              | N        |        |
|         |         |             |
 +-------------+-----+-------+----------------+----------+--------+----------
+---------+---------+-------------+

 === Fields ===
 +----+----------------+--------------+----------+---------+----------
+------------+----------+
 | id |      name      |     type     | nullable | default | shardKey |
primaryKey | identity |
 +----+----------------+--------------+----------+---------+----------
+------------+----------+
 |  1 | acct_id        | Integer      | N        | NULL    | Y        |
Y          |          |
 +----+----------------+--------------+----------+---------+----------
+------------+----------+
 |  2 | profile_name   | String       | Y        | NULL    |
|          |          |
 +----+----------------+--------------+----------+---------+----------
+------------+----------+
 |  3 | account_expiry | Timestamp(9) | Y        | NULL    |
|          |          |
 +----+----------------+--------------+----------+---------+----------
+------------+----------+
 |  4 | acct_data      | Json         | Y        | NULL    |
|          |          |
 +----+----------------+--------------+----------+---------+----------
+------------+----------+


---+----------+
```

Example 2: Describe a table and display the output as JSON

```
DESC AS JSON TABLE BaggageInfo
```

Output:

```
{
  "json_version" : 1,
  "type" : "table",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
  "name" : "BaggageInfo",
  "fields" : [{
    "name" : "ticketNo",
    "type" : "LONG",
    "nullable" : false
  }, {
    "name" : "fullName",
    "type" : "STRING",
    "nullable" : true
  }, {
    "name" : "gender",
    "type" : "STRING",
    "nullable" : true
  }, {
    "name" : "contactPhone",
    "type" : "STRING",
    "nullable" : true
  }, {
    "name" : "confNo",
    "type" : "STRING",
    "nullable" : true
  }, {
    "name" : "bagInfo",
    "type" : "JSON",
    "nullable" : true
  }],
  "primaryKey" : ["ticketNo"],
  "shardKey" : ["ticketNo"]
}
```

Example 3: Describe one particular field of a table

```
DESCRIBE TABLE BaggageInfo (ticketNo)
```

Output:

```
 +----+----------+------+----------+---------+----------+------------
+----------+
 | id |   name   | type | nullable | default | shardKey | primaryKey |
identity |
 +----+----------+------+----------+---------+----------+------------
+----------+
 |  1 | ticketNo | Long | N        | NULL    | Y        | Y
|          |
 +----+----------+------+----------+---------+----------+------------
+----------+
```

Sample data to run queries

Table 1: Airline baggage tracking application

If you want to follow along with the examples, download the script baggageschema_loaddata.sql and run it as shown below. This script creates the table used in the example and loads data into the table. One sample row is shown below.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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"
      } ]
    }, {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
      "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"
  } ]
```

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
```

The `baggageschema_loaddata.sql` contains the following:

```
### Begin Script###
load -file baggageInfo.ddl
import -table baggageInfo -file baggageData.json
### End Script ###
```

Using the `load` command, run the script.

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

Table 2: Streaming Media Service - Persistent user profile store

Download the script acctstream_loaddata.sql and run it as shown below. This script creates the table used in the example and loads data into the table. One sample row 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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": [
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                {
                   "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,
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                    "date" : "2022-03-07"
                 },
                 {
                    "episodeID": 30,
               "episodeName" : "Merci",
                    "lengthMin": 42,
                    "minWatched": 42,
                    "date" : "2022-03-08"
                 }
              ]
           }
        ]
     }
  ]
}
```

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
```

The `acctstream_loaddata.sql` contains the following:

```
### Begin Script###
load -file acctstream.ddl
import -table stream_acct -file acctstreamData.json
### End Script ###
```

Using the `load` command, run the script.

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

Table 3: JSON collection table - Shopping application

The following code inserts data into the shopping application table.

The table used in shopping application is `storeAcct`. 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 *;
```

Table Hierarchies

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

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 (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 point to the name of the parent.

Characteristics of parent-child tables:

Copyright © 2022, 2025, Oracle and/or its affiliates.

- A child table inherits the primary key columns of its parent table.
- All tables in the hierarchy have the same shard key columns, which are specified in the create table statement of the root table.
- A parent table cannot be dropped before its children are dropped.
- A referential integrity constraint is not enforced in a parent-child table.

You should consider using child tables when some form of data normalization is required. Child tables can also be a good choice when modeling 1 to N relationships and also provide ACID transaction semantics when writing multiple records in a parent-child hierarchy.

About Oracle NoSQL Database SDK drivers

Learn about Oracle NoSQL Database SDK drivers.

Oracle NoSQL Database supports many of the most popular programming languages and frameworks with idiomatic language APIs and data structures, giving your application language native access to data stored in NoSQL Database. It currently supports the following programming languages and frameworks: Java, Python, Node.js(JavaScript/TypeScript), Golang, C#/.NET, and Rust.

- Java
- Python
- Go
- Node.js
- C#
- Rust

Java

Make sure that a recent version of the java jdk is installed locally on your computer.

Make sure you have `maven` installed. See Installing Maven for details. The Oracle NoSQL Database SDK for Java is available in Maven Central repository, details available here. The main location of the project is in GitHub.

You can get all the required files for running the SDK with the following POM file dependencies.

Note

Please replace the placeholder for the version of the Oracle NoSQL Java SDK in the `pom.xml` file with the exact SDK version number.

```
<dependency>
  <groupId>com.oracle.nosql.sdk</groupId>
  <artifactId>nosqldriver</artifactId>
  <version><NOSQL_JAVASDK_VERSION></version>
</dependency>
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

The Oracle NoSQL Database SDK for Java provides you with all the Java classes, methods, interfaces and examples. Documentation is available as javadoc in GitHub or from Java API Reference Guide.

Python

Make sure that python is installed in your system. You can install the Python SDK through the Python Package Index with the command given below.

```
pip3 install borneo
```

If you are using the Oracle NoSQL Database Cloud Service you will also need to install the oci package:

```
pip3 install oci
```

The main location of the project is in GitHub.The Oracle NoSQL SDK for Python provides you with all the Python classes, methods, interfaces and examples. Documentation is available in Python API Reference Guide.

Go

Make sure you have Go installed in your computer.

The Go SDK for Oracle NoSQL Database is published as a Go module. It is recommended to use the Go modules to manage dependencies for your application. Using Go modules, you don't need to download the Go SDK explicitly. Add import statements for the SDK packages to your application code as needed. For example:

```
import "github.com/oracle/nosql-go-sdk/nosqldb"
```

When you build or test your application, the build commands will automatically add new dependencies as needed to satisfy imports, updating *go.mod* and downloading the new dependencies.

The main location of the project is in GitHub. Access the online godoc for information on using the SDK and to reference Go driver packages, types, and methods.

Node.js

Download and install Node.js from Node.js Downloads. Ensure that Node Package Manager (npm) is installed along with Node.js. Install the node SDK for Oracle NoSQL Database using one of the commands shown below.

To install as a dependency of your project:

```
npm install oracle-nosqldb
```

The npm will create a `node_modules` directory in the current directory and install it there.

To install globally:

```
npm install -g oracle-nosqldb
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

The main location of the project is in GitHub. Access the Node.js API Reference Guide to reference Node.js classes, events, and global objects.

If you are using TypeScript, use `npm` to install a supported version. Use the following command to install a specific version of the Typescript.

```
npm install typescript
```

For additional information on TypeScript, see TypeScript Modules.

About the code samples:

You can use the given code samples in TypeScript or JavaScript if using the ES6 modules.

With Oracle NoSQL Database, use JavaScript with either CommonJS or ES6 modules. In each module, how you import the NoSQLClient class and other classes/types from the Node SDK varies.

- If you want to use JavaScript with CommonJS modules, import the classes/types using the '*require*' syntax. For more information, see Node.js CommonJS Modules. For example:

  ```
  const NoSQLClient = require('oracle-nosqldb').NoSQLClient;
  ```

- If you want to use JavaScript with ES6 modules, import the classes/types using the '*import*' syntax. For more information, see Node.js ECMAScript Modules. For example:

  ```
  import { NoSQLClient } from 'oracle-nosqldb';
  ```

C#

Make sure you have .NET installed in your system.

You can install the SDK from NuGet Package Manager either by adding it as a reference to your project or independently.

- Add the SDK as a Project Reference: Run the following command to create your project directory.

  ```
  dotnet newconsole -o HelloWorld
  ```

  You may add the SDK NuGet Package as a reference to your project by using .Net CLI.

  ```
  cd <your-project-directory>
  dotnet add package Oracle.NoSQL.SDK
  ```

  Alternatively, you may perform the same using NuGet Package Manager in Visual Studio.

- Independent Install: You may install the SDK independently into a directory of your choice by using nuget.exe CLI.

  ```
  nuget.exe install Oracle.NoSQL.SDK -OutputDirectory
  <your-packages-directory>
  ```

The main location of the project is in GitHub. See Oracle NoSQL Dotnet SDK API Reference for more details of all classes and methods.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Rust

Make sure the Rust binary is installed in your system. You can download and install a Rust binary release suitable for your system. The Rust SDK for Oracle NoSQL Database is published as a Rust crate. It is recommended to use the crates.io standard Rust mechanism for usage of this crate.

Add the following dependency to your `Cargo.toml` file:

```
[dependencies]
oracle-nosql-rust-sdk = "0.1.1"
```

The main location of the project is in the GitHub. The Oracle NoSQL Database SDK for Rust provides you with all the classes, methods, interfaces, and examples and is available in Rust API Reference Guide.

Obtaining a NoSQL Handle

Learn how to access tables using Oracle NoSQL Database Drivers. Start developing your application by creating a NoSQL Handle. Use the NoSQLHandle to access the tables and execute all operations.

Non-secure data store

In your application, create a NoSQLHandle which will be your connection to the Oracle NoSQL Database Proxy. Using this NoSQLHandle you could access the Oracle NoSQL Database tables and execute Oracle NoSQL Database operations.

- Java
- Python
- Go
- Node.js
- C#
- Rust

Java

The `NoSQLHandleConfig` class allows an application to specify the security configuration information which is to be used by the handle. For non-secure access, create an instance of the `StoreAccessTokenProvider` class with the no-argument constructor. Provide the reference of `StoreAccessTokenProvider` class to the `NoSQLHandleConfig` class to establish the appropriate connection.

The following is an example of creating `NoSQLHandle` that connects to a non-secure proxy.

```
// Service URL of the proxy
String endpoint = "http://<proxy_host>:<proxy_http_port>";
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
// Create a default StoreAccessTokenProvider for accessing the proxy
StoreAccessTokenProvider provider = new StoreAccessTokenProvider();

// Create a NoSQLHandleConfig
NoSQLHandleConfig config = new NoSQLHandleConfig(endpoint);

// Setup authorization provider using StoreAccessTokenProvider
config.setAuthorizationProvider(provider);

// Create NoSQLHandle using the information provided in the config
NoSQLHandle handle = NoSQLHandleFactory.createNoSQLHandle(config);
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the host you configured earlier.
- `proxy_http_port` is the port on which the proxy is listening for requests. This should match the http port you configured earlier.

Python

A handle is created by first creating a `borneo.NoSQLHandleConfig` instance to configure the communication endpoint, authorization information, as well as default values for handle configuration.

An example of acquiring a NoSQLHandle for a non-secure Oracle NoSQL Database:

```
from borneo import NoSQLHandle, NoSQLHandleConfig
from borneo.kv import StoreAccessTokenProvider
endpoint = 'http://<proxy_host>:<proxy_http_port>'
# Create the AuthorizationProvider for a not secure store:
ap = StoreAccessTokenProvider()
# create a configuration object
config = NoSQLHandleConfig(endpoint).set_authorization_provider(ap)
# create a handle from the configuration object
handle = NoSQLHandle(config)
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the host you configured earlier.
- `proxy_http_port` is the port on which the proxy is listening for requests. This should match the http port you configured earlier.

Go

The first step in connecting a `go` application to the data store is to create a `nosqldb.Client` handle used to send requests to the service. In this case, the `Endpoint` config parameter should point to the NoSQL proxy host and port location.

```
cfg:= nosqldb.Config{
    Mode:     "onprem",
    Endpoint: "http://<proxy_host>:<proxy_http_port>",
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
client, err:=nosqldb.NewClient(cfg)
...
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the host you configured earlier.
- `proxy_http_port` is the port on which the proxy is listening for requests. This should match the http port you configured earlier.

Node.js

The `NoSQLClient` class represents the main access point to the service. To create instance of `NoSQLClient` you need to provide appropriate configuration information.

To connect to the proxy in non-secure mode, you need to specify communication endpoint.

Use the following code to connect to the proxy.

```
import { NoSQLClient, ServiceType } from 'oracle-nosqldb';
const client = new NoSQLClient({
    serviceType: ServiceType.KVSTORE,
    endpoint: '<proxy_host>:<proxy_http_port>'
});
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the host you configured earlier.
- `proxy_http_port` is the port on which the proxy is listening for requests. This should match the http port you configured earlier.

You may also choose to store the same configuration in a file. Create file *config.json* with following contents:

```
{
    "serviceType": "KVSTORE",
    "endpoint": "<proxy_host>:<proxy_http_port>",
}
```

Then you can use this sample file to create a NoSQLClient instance:

```
import { NoSQLClient } from 'oracle-nosqldb';
const client = new NoSQLClient('</path/to/config.json>');
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

Copyright © 2022, 2025, Oracle and/or its affiliates.

C#

Class `NoSQLClient` represents the main access point to the service. To create instance of `NoSQLClient` you need to provide appropriate configuration information.

In non-secure mode, the driver communicates with the proxy via the HTTP protocol. The only information required is the communication *endpoint*. For on-premise NoSQL Database, the endpoint specifies the url of the proxy, in the form *http://proxy_host:proxy_http_port*.

You can provide an instance of NoSQLConfig either directly or in a JSON configuration file.

```
var client = new NoSQLClient(
    new NoSQLConfig
    {
        ServiceType = ServiceType.KVStore,
        Endpoint = "<proxy_host>:<proxy_http_port>"
    });
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the host you configured earlier.
- `proxy_http_port` is the port on which the proxy is listening for requests. This should match the http port you configured earlier.

You may also choose to provide the same configuration in JSON configuration file. Create file *config.json* with following contents:

```
{
    "ServiceType": "KVStore",
    "Endpoint": "<proxy_host>:<proxy_http_port>"
}
```

Then you may use this file to create `NoSQLClient` instance:

```
varclient = new NoSQLClient("</path/to/config.json>");
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

Rust

In non-secure mode, the driver communicates with the proxy via the HTTP protocol. The only information required is the communication endpoint. For on-premise NoSQL Database, the endpoint specifies the URL of the proxy, in the form *http://proxy_host:proxy_http_port*.

```
let handle = Handle::builder()
.mode(HandleMode::Onprem)?
.endpoint("http://<proxy_host>:<proxy_http_port>")?
.build().await?;
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Secure data store

In your application, create a `NoSQLHandle` to connect to the secure data store through the proxy. Using the`NoSQLHandle` you could access the Oracle NoSQL Database tables and execute Oracle NoSQL Database operations. Before you start up the proxy, you need to create a bootstrap user (`proxy_user`) in the secure data store for the proxy to bootstrap its security connection. See Create a user and start proxy for a secure data store for more details.

You also need to create an application user for your application to access the secure data store. The application user will connect to the data store and perform various database operations.

```
sql-> CREATE USER <appln_user> IDENTIFIED BY "<applnuser_password>"
```

Your application user should be given a role based on the *least privilege* access, carefully balancing the needs of the application with security concerns. See Configuring privileges and roles for more details.

- Java
- Python
- Go
- Node.js
- C#
- Rust

Java

The first step for a Java application is to create a `NoSQLHandle` which will be used to send requests to the secure data store. The handle is configured using your credentials and other authentication information.

You can connect to a secure data store using the following steps.

1. Create an application user (`appln_user`) to access the data store through the secure proxy as discussed above.
2. Install the Oracle NoSQL Database Java Driver in the application's classpath.
3. For secure access, create an instance of the `StoreAccessTokenProvider` class with the parameterized constructor, and configure the NoSQL handle to use it. Use the following code to connect to the proxy.

```
String endpoint = "https://<proxy_host>:<proxy_https_port>";
StoreAccessTokenProvider atProvider =
    new
StoreAccessTokenProvider("<appln_user>","<applnuser_password>".toCharArray(
));
NoSQLHandleConfig config = new NoSQLHandleConfig(endpoint);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
config.setAuthorizationProvider(atProvider);
NoSQLHandle handle = NoSQLHandleFactory.createNoSQLHandle(config);
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the proxy host you configured earlier.
- `proxy_https_port` is the port on which the proxy is listening for requests. This should match the proxy https port configured earlier.
- `appln_user` is the user created to connect to the secure store. This should match the user created in the above section.
- `applnuser_password` is the password of the `appln_user`.

4. You can specify the details of the trust store containing the SSL certificate for the proxy in one of the following two ways.
You can set it as part of your Java code as shown below:

```
/* the trust store containing SSL cert for the proxy */
System.setProperty("javax.net.ssl.trustStore", trustStore);
if (trustStorePassword != null) {

System.setProperty("javax.net.ssl.trustStorePassword",trustStorePassword);
}
```

Alternatively, you can start-up the application program and set the `driver.trust` file's path to the Java trust store by using the following command. This is required as the proxy is already set up using the `certificate.pem` and `key-pkcs8.pem` files.

```
java -Djavax.net.ssl.trustStore=<fullpath_driver.trust> \
-Djavax.net.ssl.trustStorePassword=<password of driver.trust> \
-cp .:lib/nosqldriver.jar application_program
```

The `driver.trust` contains the `certificate.pem` or `rootCA.crt` certificate. If the certificate `certificate.pem` is in a chain signed by a trusted CA that is listed in `JAVA_HOME/jre/lib/security/cacerts`, then you don't need to append Java environment parameter `-Djavax.net.ssl.trustStore` in the Java command.

Python

A handle is created by first creating a `borneo.NoSQLHandleConfig` instance to configure the communication endpoint, authorization information, as well as default values for handle configuration.

You can connect to a secure data store using the following steps.

1. Create an application user (`appln_user`) to access the data store through the secure proxy as discussed above.
2. If running a secure store, a certificate path should be specified through the `REQUESTS_CA_BUNDLE` environment variable:

```
$ export REQUESTS_CA_BUNDLE=
<full-qualified-path-to-certificate/certificate.pem:$REQUESTS_CA_BUNDLE
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

or `borneo.NoSQLHandleConfig.set_ssl_ca_certs()`.

3. Use the following code to connect to the proxy.

```
from borneo import NoSQLHandle, NoSQLHandleConfig
from borneo.kv import StoreAccessTokenProvider
endpoint = 'https://<proxy_host>:<proxy_https_port>'
# Create the AuthorizationProvider for a secure store:
ap = StoreAccessTokenProvider(<appln_user>, <applnuser_password>)
# create a configuration object
config = NoSQLHandleConfig(endpoint).set_authorization_provider(ap)
# set the certificate path if running a secure store
config.set_ssl_ca_certs(<ca_certs>)
# create a handle from the configuration object
handle = NoSQLHandle(config)
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the proxy host you configured earlier.
- `proxy_https_port` is the port on which the proxy is listening for requests. This should match the proxy https port configured earlier.
- `appln_user` is the user created to connect to the secure store. This should match the user created in the above section.
- `applnuser_password` is the password of the `appln_user`.

Go

The first step in Oracle NoSQL Database `go` application is to create a `nosqldb.Client` handle used to send requests to the service. The handle is configured using your credentials and other authentication information:

You can connect to a secure data store using the following steps.

1. Create an application user (`appln_user`) to access the data store through the secure proxy as discussed above.
2. Use the following code to connect to the proxy.

```
import (
    "fmt"
    "github.com/oracle/nosql-go-sdk/nosqldb"
    "github.com/oracle/nosql-go-sdk/nosqldb/httputil"
)
...cfg:= nosqldb.Config{
    Endpoint: "https://<proxy_host>:<proxy_https_port>",
    Mode:     "onprem",
    Username: "<appln_user>",
    Password: "<applnuser_password>",
   },
    // Specify the CertPath and ServerName
    // ServerName is used to verify the hostname for self-signed
certificates.
    // This field is set to the "CN" subject value from the certificate
specified by CertPath.
    HTTPConfig: httputil.HTTPConfig{
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        CertPath: "<fully_qualified_path_to_cert>",
        ServerName: "<server_name>",
    },
}
client, err:=nosqldb.NewClient(cfg)
iferr!=nil {
    fmt.Printf("failed to create a NoSQL client: %v\n", err)
    return
}
deferclient.Close()
// Perform database operations using client APIs.
// ...
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the proxy host you configured earlier.
- `proxy_https_port` is the port on which the proxy is listening for requests. This should match the proxy https port configured earlier.
- `appln_user` is the user created to connect to the secure store. This should match the user created in the above section.
- `applnuser_password` is the password of the `appln_user`.

Node.js

To create instance of NoSQLClient you need to provide appropriate configuration information. You can connect to a secure data store using the following steps.

1. Create an application user (`appln_user`) to access the data store through the secure proxy as discussed above.
2. In secure mode the proxy requires the SSL Certificate and Private key If the root certificate authority (CA) for your proxy certificate is not one of the trusted root CAs , the driver needs the certificate chain file (e.g. *certificates.pem*) or a root CA certificate file (e.g. *rootCA.crt*) in order to connect to the proxy. If you are using self-signed certificate instead, the driver will need the certificate file (e.g. *certificate.pem*) for the self-signed certificate in order to connect.
   To provide the certificate or certificate chain to the driver, you have two options , either specifying in the code or setting as environment variables.

   You can specify the certificates through `httpOpt` property while creating the NoSQL handle. Inside `httpOpt` you can use `ca` property to specify the CA as shown below.

   ```
   const client = new NoSQLClient({ .....,
     httpOpt: {
       ca: fs.readFileSync(<caCertFile>)
     },.....
   });
   ```

   > **Note**
   >
   > If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Alternatively, before running your application, set the environment variable `NODE_EXTRA_CA_CERTS` as shown below.

```
export NODE_EXTRA_CA_CERTS="<fully_qualified_path_to_driver.trust>"
```

where *driver.trust* is either a certificate chain file (*certificates.pem*) for your CA, your root CA's certificate (*rootCA.crt*) or a self-signed certificate (*certificate.pem*).

3. To connect to the proxy in secure mode, in addition to communication endpoint, you need to specify user name and password of the driver user. This information is passed in `Config#auth` object under `kvstore` property and can be specified in one of 3 ways as described below.
You may choose to specify user name and password directly:

```
const NoSQLClient = require('oracle-nosqldb').NoSQLClient;
const client = new NoSQLClient({
    endpoint: 'https://<proxy_host>:<proxy_https_port>',
    auth: {
        kvstore: {
            user: '<appln_user>',
            password: '<applnuser_password>'
        }
    }
});
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the proxy host you configured earlier.
- `proxy_https_port` is the port on which the proxy is listening for requests. This should match the proxy https port configured earlier.
- `appln_user` is the user created to connect to the secure store. This should match the user created in the above section.
- `applnuser_password` is the password of the `appln_user`.

This option is less secure because the password is stored in plain text in memory.

You may choose to store credentials in a separate file which is protected by file system permissions, thus making it more secure than previous option, because the credentials will not be stored in memory, but will be accessed from this file only when login is needed. Credentials file should have the following format:

```
{
    "user":     "<appln_user>",
    "password": "<applnuser_password>"
}
```

Then you may reference this credentials file as following:

```
const NoSQLClient = require('oracle-nosqldb').NoSQLClient;
const client = new NoSQLClient({
    endpoint: 'https://<proxy_host>:<proxy_https_port>',
    auth: {
        kvstore: {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                credentials: '<path/to/credentials.json>'
            }
        }
});
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

You may also reference `credentials.json` in the configuration file used to create `NoSQLClient` instance.

```
{
    "endpoint": "https://<proxy_host>:<proxy_https_port>",
    "auth": {
        "kvstore": {
            "credentials": "<path/to/credentials.json>"
        }
    }
}


const NoSQLClient = require('oracle-nosqldb').NoSQLClient;
const client = new NoSQLClient('</path/to/config.json>');
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

C#

To create instance of `NoSQLClient` you need to provide appropriate configuration information. You can connect to a secure data store using the following steps.

1. Create an application user (`appln_user`) to access the data store through the secure proxy as discussed above.
2. To connect to the proxy in secure mode, in addition to communication endpoint, you need to specify user name and password of the driver user. This information is passed in the instance of `KVStoreAuthorizationProvider` and can be specified in any of the ways as described below.
   You may choose to specify user name and password directly:

```
var client = new NoSQLClient(
    new NoSQLConfig
    {
        Endpoint = "https://<proxy_host>:<proxy_https_port>",
        AuthorizationProvider = new KVStoreAuthorizationProvider(
            <appln_user>, // user name as string
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            <applnuser_password>) // password as char[]
    });
```

where,

- `proxy_host` is the hostname of the machine running the proxy service. This should match the proxy host you configured earlier.
- `proxy_https_port` is the port on which the proxy is listening for requests. This should match the proxy https port configured earlier.
- `appln_user` is the user created to connect to the secure store. This should match the user created in the above section.
- `applnuser_password` is the password of the `appln_user`.

This option is less secure because the password is stored in plain text in memory for the lifetime of `NoSQLClient` instance. Note that the password is specified as char[] which allows you to erase it after you are finished using NoSQLClient.

You may choose to store credentials in a separate file which is protected by file system permissions, thus making it more secure than the previous option, because the credentials will not be stored in memory, but will be accessed from this file only when the login to the store is required. Credentials file should have the following format:

```
{
    "UserName": "<appln_user>",
    "Password": "<applnuser_password>"
}
```

Then you may use this credentials file as following:

```
var client = new NoSQLClient(
    new NoSQLConfig
    {
        Endpoint: 'https://<proxy_host>:<proxy_https_port>',
        AuthorizationProvider = new KVStoreAuthorizationProvider(
            "<path/to/credentials.json>")
    });
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

You may also reference credentials.json in the JSON configuration file used to create `NoSQLClient` instance:

```
{
    "Endpoint": "https://<proxy_host>:<proxy_https_port>",
    "AuthorizationProvider": {
        "AuthorizationType": "KVStore",
        "CredentialsFile": "<path/to/credentials.json>"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        }
    }


    var client = new NoSQLClient("</path/to/config.json>");
```

Note that in `config.json` the authorization provider is represented as a JSON object with the properties for `KVStoreAuthorizationProvider` and an additional AuthorizationType property indicating the type of the authorization provider, which is KVStore for the secure on-premises store.

You need to provide trusted root certificate to the driver if the certificate chain for your proxy certificate is not rooted in one of the well known CAs. The provided certificate may be either your custom CA or self-signed proxy certificate. It must be specified using TrustedRootCertificateFile property, which sets a file path (absolute or relative) to a PEM file containing one or more trusted root certificates (multiple roots are allowed in this file). This property is specified as part of ConnectionOptions in NoSQLConfig.

```
var client = new NoSQLClient(
    new NoSQLConfig {
       Endpoint: 'https://<proxy_host>:<proxy_https_port>',
       AuthorizationProvider = new KVStoreAuthorizationProvider( "<path/to/
credentials.json>"),
       ConnectionOptions: { "TrustedRootCertificateFile": "<path/to/
certificates.pem>" }
});
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

Rust

You can connect to a secure data store using the following steps.

1. Create an application user (`appln_user`) to access the data store through the secure proxy as discussed above.
2. To connect to the proxy in secure mode, in addition to communication endpoint, you need to specify the user name and password of the driver user. This information is passed as shown below.

```
    {
        let handle = Handle::builder()
        .endpoint("https://<proxy_host>:<proxy_https_port>")?
        .mode(HandleMode::Onprem)?
        .onprem_auth("<appln_user>", "<applnuser_password>")?
        .build().await?;
    }
```

where

Copyright © 2022, 2025, Oracle and/or its affiliates.

- `proxy_host` is the host name of the machine running the proxy service. This should match the proxy host you configured earlier.
- `proxy_https_port` is the port on which the proxy is listening for requests. This should match the proxy https port configured earlier.
- `appln_user` is the user created to connect to the secure store. This should match the user created in the above section.
- `applnuser_password` is the password of the `appln_user`.

Instead of providing the credentials directly, you can also store the credentials in a separate file which is protected by file system permissions. Credentials file should have the following format:

```
{
    "UserName": "<appln_user>",
    "Password": "<applnuser_password>"
}
```

Then you may use this credentials file as following:

```
    {
        let handle = Handle::builder()
        .endpoint("https://<proxy_host>:<proxy_https_port>")?
        .mode(HandleMode::Onprem)?
        .onprem_auth_from_file("/path/to/user_pass_file")?
        .build().await?;
    }
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

You need to provide trusted root certificate to the driver if the certificate chain for your proxy certificate is not rooted in one of the well known certificate authorities (CA). The provided certificate may be either your custom certificate authority or self-signed proxy certificate.

```
{
        let handle = Handle::builder()
        .endpoint("https://<proxy_host>:<proxy_https_port>")?
        .mode(HandleMode::Onprem)?
        .onprem_auth_from_file("/path/to/user_pass_file")?
        .add_cert_from_pemfile("/path/to/certificate.pem")?
        .build().await?;
}
```

> **Note**
>
> If a file path is supplied, the path can be absolute or relative to the current working directory of the application.

Copyright © 2022, 2025, Oracle and/or its affiliates.

If you want to instruct the client to skip verifying the server's certificate, you can specify as follows:

```
{
        let handle = Handle::builder()
        .endpoint("https://<proxy_host>:<proxy_https_port>")?
        .mode(HandleMode::Onprem)?
        .onprem_auth_from_file("/path/to/user_pass_file")?
        .danger_accept_invalid_certs(true)?
        .build().await?;
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

2
Create

The articles in this section include examples to create various database objects.

Creating a namespace

A namespace defines a group of tables, within which all of the table names must be uniquely identified. 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 can't 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.

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

```
CREATE NAMESPACE [IF NOT EXISTS] namespace_name
```

Note

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

The following statement defines a namespace named `ns1`.

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

Using APIs to create namespaces:

- Java
- Python
- Go
- Node.js
- C#

Copyright © 2022, 2025, Oracle and/or its affiliates.

Java

You can create a namespace using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation. These operations are asynchronous and completion needs to be checked.

Download the full code ***Namespaces.java*** from the examples here.

```
private static void createNS(NoSQLHandle handle) throws Exception {
   String createNSDDL = "CREATE NAMESPACE IF NOT EXISTS ns1";
   SystemRequest sysreq = new SystemRequest();
   sysreq.setStatement(createNSDDL.toCharArray());
   SystemResult sysres = handle.systemRequest(sysreq);
   sysres.waitForCompletion(handle, 60000,1000);
   System.out.println("Namespace " + nsName + " is created");
}
```

Python

You can create a namespace using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation.

Download the full code ***Namespaces.py*** from the examples here.

```
def create_ns(handle):
   statement = '''CREATE NAMESPACE IF NOT EXISTS ns1'''
   sysreq = SystemRequest().set_statement(statement)
   sys_result = handle.system_request(sysreq)
   sys_result.wait_for_completion(handle, 40000, 3000)
   print('Created namespace: ns1')
```

Go

You can create a namespace using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operations. These are potentially long-running operations and completion of the operation needs to be checked.

Download the full code ***Namespaces.go*** from the examples here.

```
func createNS(client *nosqldb.Client, err error)(){
    stmt := fmt.Sprintf("CREATE NAMESPACE IF NOT EXISTS ns1")
    sysReq := &nosqldb.SystemRequest{
        Statement: stmt,
    }
    sysRes, err := client.DoSystemRequest(sysReq)
    _, err = sysRes.WaitForCompletion(client, 60*time.Second, time.Second)
    if err != nil {
        fmt.Printf("Error finishing CREATE NAMESPACE request: %v\n", err)
        return
    }
    fmt.Println("Created Namespace ns1 ")
    return
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Node.js

You can create namespace using `adminDDL` method. The `adminDDL` method is used to perform any table-independent administrative operation.

Download the full JavaScript code ***Namespaces.js*** from here and the full TypeScript code ***Namespaces.ts*** from here.

```
async function createNS(handle) {
   const createNS = `CREATE NAMESPACE IF NOT EXISTS ns1`;
   let res = await handle.adminDDL(createNS);
   console.log('Namespace created: ns1' );
}
```

C#

The `ExecuteAdminSync` method is used to perform any table-independent administrative operations.

Download the full code ***Namespaces.cs*** from the examples here.

```
private static async Task createNS(NoSQLClient client){
   var sql =
        $@"CREATE NAMESPACE IF NOT EXISTS ns1";
   var adminResult = await client.ExecuteAdminAsync(sql);
   // Wait for the operation completion
   await adminResult.WaitForCompletionAsync();
   Console.WriteLine("  Created namespace ns1");
}
```

Creating a region

Oracle NoSQL Database supports Multi-Region architecture in which you can create tables in multiple KVStores and Oracle NoSQL Database will automatically replicate inserts, updates, and deletes in a multi-directional fashion across all regions for which the table spans. Each KVStore cluster in a Multi-Region NoSQL Database setup is called a Region.

Example 1: The following CREATE REGION statement creates a remote region named `my_region1`.

```
CREATE REGION my_region1
```

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.

Example 2: Create a table in a region.

```
CREATE TABLE stream_acct_region(acct_id INTEGER,
acct_data JSON,
PRIMARY KEY(acct_id)) IN REGIONS my_region1
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

> **Note**
>
> The region `my_region1` should be set as the local region before creating the table.

Using APIs to create regions:

- Java
- Python
- Go
- Node.js
- C#

Java

You can create a region using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation. Once the region is created, a table can be created and added to the region using the `TableRequest` class.

Download the full code ***Regions.java*** from the examples here.

```
/* Create a remote region and a local region*/
private static void crtRegion(NoSQLHandle handle) throws Exception {
    // Create a remote region
    String createRemRegDDL = "CREATE REGION "+ remRegName;
    SystemRequest sysreq1 = new SystemRequest();
    sysreq1.setStatement(createRemRegDDL.toCharArray());
    SystemResult sysres1 = handle.systemRequest(sysreq1);
    sysres1.waitForCompletion(handle, 60000,1000);
    System.out.println(" Remote Region " + remRegName + " is created");
    // Create a local region
    String createLocRegDDL = "SET LOCAL REGION "+ localRegName;
    SystemRequest sysreq2 = new SystemRequest();
    sysreq2.setStatement(createLocRegDDL.toCharArray());
    SystemResult sysres2 = handle.systemRequest(sysreq2);
    sysres2.waitForCompletion(handle, 60000,1000);
    System.out.println(" Local Region " + localRegName + " is created");
}

/**
 * Create a table and add the table in a region
 */
private static void crtTabInRegion(NoSQLHandle handle) throws Exception {
    String createTableDDL = "CREATE TABLE IF NOT EXISTS " + tableName +
                                                              "(acct_Id
INTEGER," +
                                                              "profile_name
STRING," +
                                                              "account_expiry
TIMESTAMP(1) ," +
                                                              "acct_data
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
JSON, " +
                                                                    "PRIMARY
KEY(acct_Id)) IN REGIONS FRA";

    TableRequest treq = new TableRequest().setStatement(createTableDDL);
    TableResult tres = handle.tableRequest(treq);
    /* The request is async,
     * so wait for the table to become active.
     */
    System.out.println("Table " + tableName + " is active");
}
```

Python

You can create a region using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation. After a region is created, a table can be created and added to the region using the `borneo.TableRequest` class.

Download the full code ***Regions.py*** from the examples here.

```
# create a remote and local region
def create_region(handle):
   #Create a remote region
   statement = '''CREATE REGION LON'''
   sysreq = SystemRequest().set_statement(statement)
   sys_result = handle.system_request(sysreq)
   sys_result.wait_for_completion(handle, 40000, 3000)
   print('Remote region LON is created')
   #Create a local region
   statement1 = '''SET LOCAL REGION FRA'''
   sysreq1 = SystemRequest().set_statement(statement1)
   sys_result1 = handle.system_request(sysreq1)
   sys_result1.wait_for_completion(handle, 40000, 3000)
   print('Local region FRA is created')

#Create a table in the local region
def create_tab_region(handle):
   statement = '''create table if not exists stream_acct (acct_Id INTEGER,
                                                           profile_name
STRING,
                                                           account_expiry
TIMESTAMP(1),
                                                           acct_data JSON,
                                                           primary
key(acct_Id)) IN REGIONS FRA'''
   request = TableRequest().set_statement(statement)
   # Ask the cloud service to create the table, waiting for a total of 40000
milliseconds
   # and polling the service every 3000 milliseconds to see if the table is
active
   table_result = handle.do_table_request(request, 40000, 3000)
   table_result.wait_for_completion(handle, 40000, 3000)
   if (table_result.get_state() == State.ACTIVE):
      print('Created table: stream_acct')
   else:
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        raise NameError('Table stream_acct is in an unexpected state ' +
str(table_result.get_state()))
```

Go

You can create a region using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operations. After a region is created, a table can be created and added to the region using the `TableRequest` class.

Download the full code ***Regions.go*** from the examples here.

```
//Creates a remote and a local region
func crtRegion(client *nosqldb.Client, err error)(){
   // Create a remote region
   stmt := fmt.Sprintf("CREATE REGION LON")
   sysReq := &nosqldb.SystemRequest{
    Statement: stmt,
   }
   sysRes, err := client.DoSystemRequest(sysReq)
   _, err = sysRes.WaitForCompletion(client, 60*time.Second, time.Second)
   if err != nil {
      fmt.Printf("Error finishing CREATE REGION request: %v\n", err)
      return
   }
   fmt.Println("Created REGION LON ")
   // Create a local region
   stmt1 := fmt.Sprintf("SET LOCAL REGION FRA")
   sysReq1 := &nosqldb.SystemRequest{
        Statement: stmt1,
   }
   sysRes1, err1 := client.DoSystemRequest(sysReq1)
   _, err1 = sysRes1.WaitForCompletion(client, 60*time.Second, time.Second)
   if err1 != nil {
      fmt.Printf("Error finishing CREATE REGION request: %v\n", err)
      return
   }
   fmt.Println("Created REGION FRA ")
   return
}

//creates a table in a specific region
func crtTabInRegion(client *nosqldb.Client, err error, tableName string)(){
   stmt := fmt.Sprintf("CREATE TABLE IF NOT EXISTS %s ("+
        "acct_Id INTEGER," +
        "profile_name STRING," +
        "account_expiry TIMESTAMP(1) ," +
        "acct_data JSON, " +
        "PRIMARY KEY(acct_Id)) IN REGIONS FRA",tableName)
   tableReq := &nosqldb.TableRequest{
        Statement: stmt
   }
   tableRes, err := client.DoTableRequest(tableReq)
   if err != nil {
      fmt.Printf("cannot initiate CREATE TABLE request: %v\n", err)
      return
   }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    // The create table request is asynchronous, wait for table creation to
complete.
    _, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)
    if err != nil {
       fmt.Printf("Error finishing CREATE TABLE request: %v\n", err)
       return
    }
    fmt.Println("Created table ", tableName)
    return
}
```

Node.js

You can create a region using `adminDDL` method. The `adminDDL` method is used to perform any table-independent administrative operation. After a region is created, a table can be created and added to the region using the `tableDDL` method.

Download the full JavaScript code ***Regions.js*** from the examples here and the full TypeScript code ***Regions.ts*** from the examples here.

```
//creates a remote and a local region
async function createRegion(handle) {
   // Create a remote region
   const crtRemReg = `CREATE REGION LON`;
   let res = await handle.adminDDL(crtRemReg);
   console.log('Remote region created: LON' );
   // Create a local region
   const crtLocalReg = `SET LOCAL REGION FRA`;
   let res1 = await handle.adminDDL(crtLocalReg);
   console.log('Local region created: FRA' );
}

//creates a table in a given region
async function crtTabInRegion(handle) {
   const createDDL = `CREATE TABLE IF NOT EXISTS ${TABLE_NAME} (acct_Id
INTEGER,
                                                                 profile_name
STRING,

account_expiry TIMESTAMP(1),
                                                                 acct_data
JSON,
                                                                 primary
key(acct_Id)) IN REGIONS FRA`;
   let res =  await handle.tableDDL(createDDL, {
            complete: true
   });
   console.log('Table created: ' + TABLE_NAME);
}
```

C#

You can create a region using `ExecuteAdminSync` module. The `ExecuteAdminSync` method is used to perform any table-independent administrative operations. After a region is created, a table can be created and added to the region using either `ExecuteTableDDLAsync` or `ExecuteTableDDLWithCompletionAsync` methods.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***Regions.cs*** from the examples here.

```
private static async Task createRegion(NoSQLClient client){
   // Create a remote region
   var sql = $@"CREATE REGION LON";
   var adminResult = await client.ExecuteAdminAsync(sql);
   // Wait for the operation completion
   await adminResult.WaitForCompletionAsync();
   Console.WriteLine("  Created remote REGION LON");
   // Create a local region
   var sql1 = $@"SET LOCAL REGION FRA";
   var adminResult1 = await client.ExecuteAdminAsync(sql1);
   // Wait for the operation completion
   await adminResult1.WaitForCompletionAsync();
   Console.WriteLine("  Created local REGION FRA");
}

private static async Task createTabInRegion(NoSQLClient client){
   // Create a table
   var sql =
      $@"CREATE TABLE IF NOT EXISTS {TableName}(acct_Id INTEGER,
                                          profile_name STRING,
                                          account_expiry TIMESTAMP(1),
                                          acct_data JSON,
                                          primary key(acct_Id)) IN
REGIONS FRA";
   Console.WriteLine("\nCreate table {0}", TableName);
   var tableResult = await client.ExecuteTableDDLAsync(sql);
   // Wait for the operation completion
   await tableResult.WaitForCompletionAsync();
   Console.WriteLine("  Table {0} is created",tableResult.TableName);
   Console.WriteLine("  Table state: {0}", tableResult.TableState);
}
```

Creating a table

The table is the basic structure to hold user data. You use a SQL command (`CREATE TABLE` statement) or `TableRequest` API commands to create a new table.

Guidelines for creating a table:

- The table definition must include at least one field definition and exactly one primary key definition. For more information on primary key definition, see Create Table.
- The field definition specifies the name of the column, its data type, whether the column is nullable or not, an optional default value, whether or not the column is an IDENTITY column and an optional comment. All fields (other than the PRIMARY KEY) are nullable by default.
- 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 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

automatically by the 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.

- You specify the REGIONS clause if the table being created is a Multi-Region table. The REGIONS clause lists all the regions that the table should span.

> **Note**
>
> The JSON collection tables have exactly one primary key definition.

Create a table :

- Using SQL commands
- Using TableRequest API

Using SQL commands

You can use CREATE TABLE command in SQL to create NoSQL tables.

The following section highlights different options that can be used while creating a table using the CREATE TABLE DDL statement.

Example 1: Create an airline baggage tracking application table that holds 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)
)
```

In the schema above, you use the CREATE TABLE statement to define a `BaggageInfo` table. 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 fixed schema. The `bagInfo` column is a schema-less JSON array, which represents the tracking information of a passenger's checked-in baggage.

For more details on airline baggage tracking application, see Airline baggage tracking application.

Example 2: Create a streaming media service table with a JSON field to keep track of the subscriber's current activity.

```
CREATE TABLE stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data JSON,
PRIMARY KEY(acct_id)
)USING TTL 5 DAYS
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

In the schema above, you use the CREATE TABLE statement to create a `stream_acct` table. The subscriber's account ID, `acct_id` field is the primary key in this table. The fields `profile_name`, `account_expiry` store the viewership details, which is a part of fixed schema. The `acct_data` column is a schema-less JSON field, which stores the details of the shows viewed by a subscriber.

For more details on streaming media service application, see Streaming Media Service - Persistent user profile store

You also specify a TTL value, after which the rows automatically expire and are not available anymore. The TTL value must be in 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_hours` functions. For details, see Functions on Rows.

Example 3: Create a streaming media service table with various fixed-schema definitions.

```
CREATE TABLE IF NOT EXISTS stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data RECORD (
firstName STRING,
lastName STRING,
country STRING,
shows JSON
),
PRIMARY KEY(acct_id)
)
```

In the schema above, you define a variation of the fixed-schema by including the `acct_data` field as a RECORD data type.

A record is an ordered collection of one or more key-item pairs. The keys in a record must be strings and the associated items can be of different data types. The fields in a record are a part of the fixed-schema and you will not be able to add or remove them. In the example above, the fields `firstName`, `lastName`, `country`, and `shows` are the keys for the `acct_data` record. Defining a record is helpful when you want to store data as part of a bigger data set. You can insert/update/fetch the whole subset in a record using the field step expressions.

You can also nest the records as follows:

```
CREATE TABLE IF NOT EXISTS  stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data RECORD (
   firstName STRING,
   lastName STRING,
   country STRING,
   shows RECORD (
           showName STRING,
           showId INTEGER,
           type JSON,
           numSeasons INTEGER,
           seriesInfo ARRAY(JSON)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   )
),
PRIMARY KEY(acct_id)
)
```

The `shows` field is a nested RECORD type used to track the details of the viewed shows.

Example 4: Create a streaming media service table as a hierarchical table structure.

In the following schemas, you create `stream_acct` table as a parent table and `acct_data` table as a child table:

```
CREATE TABLE IF NOT EXISTS stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
PRIMARY KEY(acct_id))


CREATE TABLE IF NOT EXISTS stream_acct.acct_data(
profile_id INTEGER,
user_name STRING,
firstName STRING,
lastName  STRING,
country   STRING,
shows     JSON,
PRIMARY KEY(profile_id))
```

With the parent-child table definition above, the streaming media service can support multiple user profiles under a single subscription.

You define the `acct_data` table as a child table with a primary key `profile_id` to identify a user's profile. In addition to defining a primary key for the table, the `acct_data` table implicitly inherits the `acct_id` primary key column of its parent `stream_acct` table.

You can define multiple child tables under the same `stream_acct` parent table. You can further define child tables for the `acct_data` table. All the tables in the hierarchy have the same shard key column, which is specified in the create table statement of the highest parent table in the hierarchy. In this example, the primary key `acct_id` of the parent table is also the shard key for the `stream_acct` and `acct_data` tables.

Example 5: Create a streaming media service table with primary key as an IDENTITY column.

```
CREATE TABLE IF NOT EXISTS stream_acct(
acct_id INTEGER GENERATED BY DEFAULT AS IDENTITY (START WITH 4 INCREMENT BY 1
NO CYCLE),
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data JSON,
PRIMARY KEY(acct_id)
)
```

In this example, you create the `stream_acct` table with the same fields defined in Example 2. However, the schema defines the primary key `acct_id` as an IDENTITY column. When you define an IDENTITY column as GENERATED BY DEFAULT AS IDENTITY, the system generates the IDENTITY column values automatically if you do not supply one. The system

Copyright © 2022, 2025, Oracle and/or its affiliates.

starts to generate `acct_id` values from 4 incrementing by 1. So, values for the `acct_id` column will be 4, 5, 6... and so forth up to the maximum value of the LONG data type. As the schema defines the NO CYCLE option, the system raises an exception after the maximum value as it has reached the end of the sequence generator.

Example 6: Create a streaming media service table with primary key as a UUID data type.

```
CREATE TABLE IF NOT EXISTS stream_acct(
acct_id STRING AS UUID GENERATED BY DEFAULT,
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data JSON,
PRIMARY KEY(acct_id)
)
```

In this example, you create the `stream_acct` table with the same fields defined in Example 2. However, the schema defines the primary key `acct_id` as a UUID data type, GENERATED BY DEFAULT. The system automatically generates a value for the UUID column if you do not supply one. The potential advantage of defining primary key as a UUID data type is while the system guarantees the uniqueness of the IDENTITY columns only within a NoSQL data store in a region, UUID data type generates a globally unique identifier for the records in a table that span multiple regions.

MR_COUNTER Data Type

Example 7: Create a streaming media service table with an MR_COUNTER data type.

```
CREATE TABLE IF NOT EXISTS  stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
firstName STRING,
lastName STRING,
country STRING,
shows JSON (counter as INTEGER MR_COUNTER),
PRIMARY KEY(acct_id)) IN REGIONS DEN,LON
```

In the schema above, you define the streaming media service table with an additional MR_COUNTER data type in the `shows` field. The MR_COUNTER data type is a Conflict-free Replicated Data Type (CRDT) counter. CRDTs provide a way for concurrent modifications to be merged across regions without user intervention. In a multi-region setup of an Oracle NoSQL Database, copies of the same data must be stored in multiple regions and data may be concurrently modified in different regions. The MR_COUNTER data type ensures that when data modifications happen simultaneously on different regions, data always gets automatically merged into a consistent state.

In this example, the system creates the table in two regions DEN and LON. Each counter field in the JSON document will track the latest count of the shows the user streams in that region. As all the Multi-Region tables in the participating regions are synchronized, the system automatically merges these concurrent modifications to reflect the latest updates of the `counter` without any user intervention.

JSON Collection Tables

Example 1: Create a streaming media service table as a JSON collection table.

Copyright © 2022, 2025, Oracle and/or its affiliates.

In the following CREATE TABLE statement, you create the `stream_acct` table as a JSON collection table:

```
CREATE TABLE IF NOT EXISTS stream_acct(
acct_id INTEGER,
PRIMARY KEY(acct_id)) AS JSON COLLECTION
```

This JSON collection table includes the `acct_id` 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 this table, the JSON collection table automatically considers the inserted fields other than the `acct_id` field to be JSON fields. You can use this JSON collection table to store and retrieve TV streaming data purely as documents.

Example 2: Create a shopping application table as a JSON collection table.

The following CREATE TABLE statement defines a `storeAcct` table, which is a JSON collection table created for a shopping application. This table includes the `contactPhone` as the primary key field of the type string.

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

To insert data into the tables, see Inserting, Modifying, and Deleting Data.

Enabling Before-images During Table Creation

The before-image of any write is the table row before it gets updated or deleted by a DML operation. This feature is particularly useful in streams subscriptions wherein, streaming before-images enables applications to compute the change or delta made by write operations to tables in a data store.

The ENABLE BEFORE IMAGE clause enables before-images generation and persistent storage of those images. For more details on using before-images in a subscription stream, see Streams Developer's Guide.

Example: Enable before-images while creating a streaming media service table.

In the following CREATE TABLE statement, you create a `stream_acct` table with before-images enabled.

```
CREATE TABLE stream_acct(
acct_id INTEGER,
profile_name STRING,
account_expiry TIMESTAMP(9),
acct_data JSON,
PRIMARY KEY(acct_id)
)USING TTL 5 DAYS ENABLE BEFORE IMAGE USING TTL 48 HOURS
```

You provide both table TTL and before-images TTL together in the CREATE TABLE statement above. Both TTL values work independently of each other.

The table TTL is set to 5 days and the before-images TTL is set to 48 hours. The generated before-images on the `stream_acct` table are stored on the disk for 48 hours. After this duration, the before-images expire freeing up the disk space and will not appear in the stream.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Without the TTL definition for before-images, the generated before-images remain for 24 hours, unless adjusted by a TTL value.

To enable before-images generation after the table is created, you must use the ALTER TABLE statement with an `ENABLE BEFORE IMAGE` clause.

Using TableRequest API

You can use TableRequest API to create NoSQL tables.

- Java
- Python
- Go
- Node.js
- C#

Java

The `TableRequest` class is used to create tables. Execution of operations specified by this request is asynchronous. These are potentially long-running operations. `TableResult` is returned from `TableRequest` operations and it encapsulates the state of the table. See Oracle NoSQL Java SDK API Reference for more details on the TableRequest class and its methods.

Download the full code ***CreateTable.java*** from the examples here.

```
private static void createTab(NoSQLHandle handle) throws Exception {
   String createTableDDL =
   "CREATE TABLE IF NOT EXISTS " + tableName +
                              "(acct_Id INTEGER," +
                              "profile_name STRING," +
                              "account_expiry TIMESTAMP(1) ," +
                              "acct_data JSON, " +
                              "PRIMARY KEY(acct_Id))";

   TableLimits limits = new TableLimits(20, 20, 1);
   TableRequest treq = new TableRequest()
          .setStatement(createTableDDL)
          .setTableLimits(limits);
   TableResult tres = handle.tableRequest(treq);
   /* The request is async,
    * so wait for the table to become active.
    */
    tres.waitForCompletion(handle, 60000,1000);
    System.out.println("Created Table: " + tableName);
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Note

Table limits are applicable for Oracle NoSQL Database Cloud Service only. If limits are set for an on-premises NoSQL Database they are silently ignored.

Creating a child table: You use the same `TableRequest` class and methods to execute DDL statement to create a child table.

While creating a child table :

- You need to specify the full name of the table (name_parent_table.name_child_table)
- Table limits need not be explicitly set as a child table inherits the limits of a parent table.

Download the full code ***TableJoins.java*** from the examples to understand how to create a parent-child table here.

Creating a JSON collection table: The JSON collection table includes documents with one or more primary key fields and JSON fields. Create a JSON collection table as follows:

```
/* Create a JSON collection table with an integer primary key*/

private static void createTable(NoSQLHandle handle) throws Exception {

String createTableDDL = "CREATE TABLE IF NOT EXISTS " + usersJSON + "(id
INTEGER," + "PRIMARY KEY(id)) AS JSON COLLECTION";

  TableRequest treq = new TableRequest().setStatement(createTableDDL);

  System.out.println("Creating table");
  TableResult tres = handle.tableRequest(treq);

/* The table request is asynchronous, so wait for the table to become
active.*/

  TableResult.waitForState(handle, tres.getTableName(),
TableResult.State.ACTIVE, 60000, 1000);
}
```

Python

The `borneo.TableRequest` class is used to create a table. All calls to `borneo.NoSQLHandle.table_request()` are asynchronous so it is necessary to check the result and call `borneo.TableResult.wait_for_completion()` to wait for the operation to complete. See Oracle NoSQL Python SDK API Reference for more details on `table_request` and its methods.

Download the full code ***CreateTable.py*** from the examples here.

```
def create_table(handle):
  statement = '''create table if not exists
              stream_acct (acct_Id INTEGER,
                           profile_name STRING,
                           account_expiry TIMESTAMP(1),
                           acct_data JSON,
                           primary key(acct_Id))'''
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    request = TableRequest().set_statement(statement)
                            .set_table_limits(TableLimits(20, 10, 1))

    table_result = handle.do_table_request(request, 40000, 3000)
    table_result.wait_for_completion(handle, 40000, 3000)

    if (table_result.get_state() == State.ACTIVE):
        print('Created table: stream_acct')
    else:
        raise NameError('Table stream_acct is in an unexpected state ' +
                         str(table_result.get_state()))
```

> **Note**
>
> Table limits are applicable for Oracle NoSQL Database Cloud Service only. If limits are set for an on-premises NoSQL Database they are silently ignored.

Creating a child table: You use the same `TableRequest` class and methods to execute DDL statement to create a child table.

While creating a child table :

- You need to specify the full name of the table (name_parent_table.name_child_table).
- Table limits need not be explicitly set as a child table inherits the limits of a parent table.

Download the full code ***TableJoins.py*** from the examples here.

Creating a JSON collection table: The JSON collection table includes documents with one or more primary key fields and JSON fields. Create a JSON collection table as follows:

```
/* Create a JSON collection table with an integer primary key */
statement = 'create table if not exists usersJSON(id integer,' + 'primary
key(id)) AS JSON COLLECTION'
print('Creating table: ' +
statement)

request = TableRequest().set_statement(statement)


/* assume that a handle has been created, as the handle, make the request */
/* wait for 60 seconds, polling every 1 seconds */
result = handle.do_table_request(request, 60000, 1000)
result.wait_for_completion(handle, 60000, 1000)
```

Go

The `TableRequest` class is used to create a table. Execution of operations specified by `TableRequest` is asynchronous. These are potentially long-running operations. This request is used as the input of a `Client.DoTableRequest()` operation, which returns a `TableResult` that can be used to poll until the table reaches the desired state. See Oracle NoSQL Go SDK API Reference for more details on the various methods of the `TableRequest` class.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***CreateTable.go*** from the examples here.

```
func createTable(client *nosqldb.Client, err error, tableName string)(){
// Creates a table
 stmt := fmt.Sprintf("CREATE TABLE IF NOT EXISTS %s ("+
         "acct_Id INTEGER," +
         "profile_name STRING," +
         "account_expiry TIMESTAMP(1) ," +
         "acct_data JSON, " +
         "PRIMARY KEY(acct_Id))",tableName)

 tableReq := &nosqldb.TableRequest{
        Statement: stmt,
        TableLimits: &nosqldb.TableLimits{
            ReadUnits:  20,
            WriteUnits: 20,
            StorageGB:  1,
              },
 }

 tableRes, err := client.DoTableRequest(tableReq)
 if err != nil {
    fmt.Printf("cannot initiate CREATE TABLE request: %v\n", err)
    return
 }
 // The create table request is asynchronous,
 // wait for table creation to complete.
 _, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)

 if err != nil {
    fmt.Printf("Error finishing CREATE TABLE request: %v\n", err)
    return
 }
 fmt.Println("Created table: ", tableName)
 return
}
```

> **Note**
>
> Table limits are applicable for Oracle NoSQL Database Cloud Service only. If limits are set for an on-premises NoSQL Database they are silently ignored.

Creating a child table: You use the same `TableRequest` class and methods to execute DDL statement to create a child table.

While creating a child table :

- You need to specify the full name of the table (name_parent_table.name_child_table).
- Table limits need not be explicitly set as a child table inherits the limits of a parent table.

Download the full code ***TableJoins.go*** from the examples here.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Creating a JSON collection table: The JSON collection table includes documents with one or more primary key fields and JSON fields. Create a JSON collection table as follows:

```
/* Create a JSON collection table with an integer primary key with a TTL of 3 
days*/
tableName := "usersJSON"
stmt := fmt.Sprintf("CREATE TABLE IF NOT EXISTS %s "+
    "(id integer, PRIMARY KEY(id)) "+
    "AS JSON COLLECTION USING TTL 3 DAYS", tableName)

tableReq := &nosqldb.TableRequest{
    Statement: stmt, }
tableRes, err := client.DoTableRequest(tableReq)
if err != nil {
    fmt.Printf("cannot initiate CREATE TABLE request: %v\n", err)
    return
}
_, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)
if err != nil {
    fmt.Printf("Error finishing CREATE TABLE request: %v\n", err)
    return
}
fmt.Println("Created table ", tableName)
```

Node.js

You can create a table using the `tableDDL` method. This method is asynchronous and it returns a Promise of `TableResult`. The `TableResult` is a plain JavaScript object that contains the status of the DDL operation such as its TableState, name, schema, and its TableLimits. For method details, see NoSQLClient class.

Download the full JavaScript code ***CreateTable.js*** from the examples here and the full TypeScript code ***CreateTable.ts*** from the examples here.

```
import { NoSQLClient, ServiceType } from 'oracle-nosqldb';
const client = new NoSQLClient('config.json');
const TABLE_NAME = 'stream_acct;
async function createTable(handle) {
  const createDDL = `CREATE TABLE IF NOT EXISTS
                ${TABLE_NAME} (acct_Id INTEGER,
                               profile_name STRING,
                               account_expiry TIMESTAMP(1),
                               acct_data JSON,
                               primary key(acct_Id))`;
  let res =  await handle.tableDDL(createDDL, {
             complete: true }
  );
  console.log('Created table: ' + TABLE_NAME);
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

After the above call returns, the result will reflect the final state of the operation. Alternatively, to use the complete option, substitute the code in the try-catch block above with the following code sample.

```
const createDDL = `CREATE TABLE IF NOT EXISTS
                  ${TABLE_NAME} (acct_Id INTEGER,
                                 profile_name STRING,
                                 account_expiry TIMESTAMP(1),
                                 acct_data JSON,
                                 primary key(acct_Id))`;
 let res =  await client.tableDDL(createDDL, {complete: true,});
console.log('Created table: ' + TABLE_NAME);
```

Creating a child table: You use the same `TableRequest` class and methods to execute DDL statement to create a child table.

While creating a child table :

- You need to specify the full name of the table (name_parent_table.name_child_table)
- Table limits need not be explicitly set as a child table inherits the limits of a parent table.

Download the full JavaScript code ***TableJoins.js*** from the examples here and the full TypeScript code ***TableJoins.ts*** from the examples here.

Creating a JSON collection table: The JSON collection table includes documents with one or more primary key fields and JSON fields. Create a JSON collection table as follows:

```
import { NoSQLClient, ServiceType } from 'oracle-nosqldb';
const client = new NoSQLClient('config.json');

/* Create a JSON collection table with an integer primary key */
const TABLE_NAME = 'usersJSON';
async function createTable() {

    const createDDL = `CREATE TABLE IF NOT EXISTS ${TABLE_NAME} (id INTEGER,
PRIMARY KEY(id)) AS JSON COLLECTION`;
    console.log('Create table: ' + createDDL);

        let res =  await client.tableDDL(createDDL, {
            complete: true,
        });
}
```

C#

To create a table use either of the methods `ExecuteTableDDLAsync` or `ExecuteTableDDLWithCompletionAsync`. Both these methods return `Task<TableResult>`. `TableResult` instance contains status of DDL operation such as TableState and table schema. See Oracle NoSQL Dotnet SDK API Reference for more details on these methods.

Download the full code ***CreateTable.cs*** from the examples here.

```
private static async Task createTable(NoSQLClient client){
// Create a table
  var sql =
    $@"CREATE TABLE IF NOT EXISTS
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        {TableName}(acct_Id INTEGER,
                    profile_name STRING,
                    account_expiry TIMESTAMP(1),
                    acct_data JSON,
                    primary key(acct_Id))";
  var tableResult = await client.ExecuteTableDDLAsync(sql,
    new TableDDLOptions{TableLimits = new TableLimits(20, 20, 1)});

  // Wait for the operation completion
  await tableResult.WaitForCompletionAsync();
  Console.WriteLine("  Created table: ",tableResult.TableName);
}
```

> **Note**
>
> Table limits are applicable for Oracle NoSQL Database Cloud Service only. If limits are set for an on-premises NoSQL Database they are silently ignored.

Creating a child table: You use the same `TableRequest` class and methods to execute DDL statement to create a child table.

While creating a child table :

- You need to specify the full name of the table (name_parent_table.name_child_table)
- Table limits need not be explicitly set as a child table inherits the limits of a parent table.

Download the full code ***TableJoins.cs*** from the examples here.

Creating a JSON collection table: The JSON collection table includes documents with one or more primary key fields and JSON fields. Create a JSON collection table as follows:

```
/* Create a JSON collection table with an integer primary key */

var client = new NoSQLClient("config.json");
try {
  var statement = "CREATE TABLE IF NOT EXISTS usersJSON(id INTEGER,"
    + "PRIMARY KEY(id)) AS JSON COLLECTION";

var result = await client.ExecuteTableDDLAsync(statement);
  await result.WaitForCompletionAsync();
  Console.WriteLine("Table users created.");
} catch(Exception ex) {
   // handle exceptions
}
```

Create and View 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 for is contained in the index.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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)

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Classification of Indexes
- Creating Indexes
- View Index

Classification of Indexes

Indexes can be classified based on fields, schema, entries, or a combination of them.

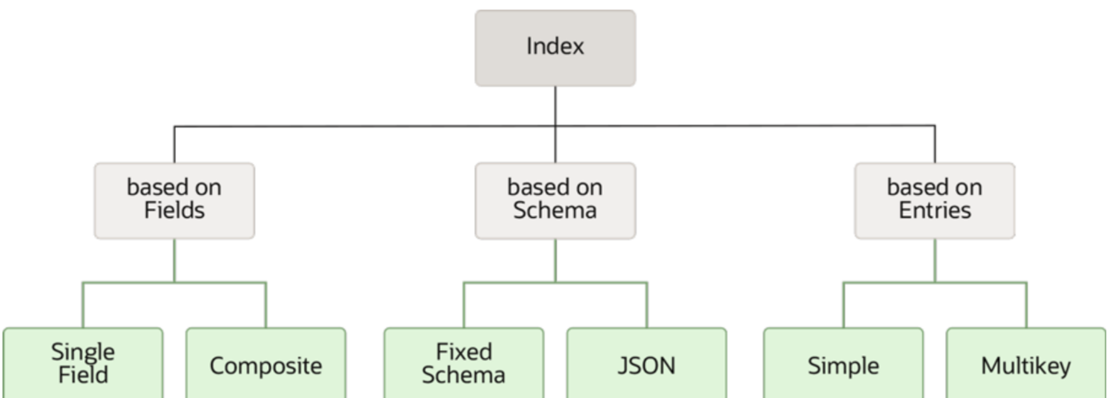

Single Field Index: An index is called a single field index if it is created on only one field of a table.

Composite Index: An index is called a composite index if it is created on more than one field of a table

Fixed Schema Index: An index is called a fixed schema index if all the fields that are indexed are strongly typed data.

> **Note**
>
> A data type is called precise if it is not one of the wild card types. Items that have precise types are said to be strongly typed.

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.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

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.

Creating Indexes

You can create an index for a NoSQL table using SQL commands or using TableRequest API.

- Using SQL commands

Copyright © 2022, 2025, Oracle and/or its affiliates.

- Using TableRequest API

Using SQL commands

An index can be created using the `CREATE INDEX` command.

Create a single field index:

Example: Create an index on passengers reservation code.

```
CREATE INDEX fixedschema_conf ON baggageInfo(confNo)
```

The above is an example of a single-column fixed schema index. The index is created on the `confNo` field having `string` data type in the `baggageInfo` table.

Create a composite index:

Example : Create an index on the full name and phone number of passengers.

```
CREATE INDEX compindex_namephone ON baggageInfo(fullName,contactPhone)
```

The above is an example of a composite index. The index is created on two fields in the baggageInfo schema, on full name and the contact phone number.

Note

You can have one or more fields of this index as fixed schema columns.

Create a 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 1: Create an index on the tag number of passengers bags.

```
CREATE INDEX jsonindex_tagnum ON baggageInfo(bagInfo[].tagnum as INTEGER)
```

The above is an example of a JSON index. The index is created on the `tagnum` field present in the `baginfo` JSON field in the `baggageInfo` table. Notice that you provide a data type for the `tagnum` field while creating the index.

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.

Example 2: Create an index on the route of passengers.

```
CREATE INDEX jsonindex_routing ON baggageInfo(bagInfo[].routing as ANYATOMIC)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Create an Index on JSON Collection Table

Indexing the fields in the JSON collection table is similar to creating JSON indexes. You must specify the name (along with the path expression) and ANYATOMIC for the type definition, or, for strongly typed indexes, you can specify the JSON type of the fields being indexed.

If you are indexing a top-level JSON field in the document, the field name is the 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. It is recommended to use ANYATOMIC in the JSON collection tables for more flexibility.

Example : Create a composite index on the JSON collection table created for a shopping application.

```
CREATE INDEX idx_ntfy_cty on storeAcct (address.city as ANYATOMIC, notify as
ANYATOMIC)
```

In the composite index above, the `notify` field is a top-level field of the `storeAcct` table, which can be indexed by specifying the field name as the path. The `city` field is nested in the `address` field and must be indexed using the path expression.

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.

Create a 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 return an array or map or a nested data type.

Example: Create an index on three fields, when the bag was last seen, the last seen station and the arrival date and time.

```
CREATE INDEX simpleindex_arrival ON baggageInfo(bagInfo[].lastSeenTimeGmt as
ANYATOMIC,
bagInfo[].bagArrivalDate as ANYATOMIC, bagInfo[].lastSeenTimeStation as
ANYATOMIC)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

The above is an example of a simple index created on a JSON document in a JSON field. The index is created on the `lastSeenTimeGmt` and `bagArrivalDate` and `lastSeenTimeStation`, all from the `bagInfo` JSON document in the info JSON field in the `baggageInfo` table. 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 `lastSeenTimeGmt` or `bagArrivalDate` or `lastSeenTimeStation` entry in the `bagInfo` JSON document, or if there is no `bagInfo` JSON array, then the special value EMPTY is indexed.

Create a 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 an array or a nested data type. In a multikey index, for each table row, index entries are created on all the elements in arrays that are being indexed.

Example 1: Multikey index: Create an index on the series info array of the streaming account application.

```
CREATE INDEX multikeyindex1 ON stream_acct
(acct_data.contentStreamed[].seriesInfo[] AS ANYATOMIC)
```

The index is created on the `seriesInfo[]` array in the stream_acct table. Here, all the elements in the `seriesInfo[]` array in each row of the `stream_acct` table will be indexed.

Example 2: Nested multikey index: Create an index on the episode details array of the streaming account application.

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.

```
CREATE INDEX multikeyindex2 ON stream_acct (
    acct_data.contentStreamed[].seriesInfo[].episodes[]  AS ANYATOMIC)
```

The above 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 `episodes[]` array in the `seriesInfo[]` array in the `acct_data` JSON of the `stream_acct` table.

Example 3: 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.

```
CREATE INDEX multikeyindex3 ON stream_acct (acct_data.country AS ANYATOMIC,
acct_data.contentStreamed[].seriesInfo[].episodes[]  AS ANYATOMIC)
```

The above is an example of a composite multikey index having one multikey index path and one simple index path. The index is created on the `country` field and `episodes[]`array in the `acct_data` JSON column of the `stream_acct` table.

See Specifications & Restrictions on Multikey index to learn about restrictions on multikey index.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Create an index with NO NULLS clause

You can create an index with the optional WITH NO NULLS clause. In that case, the rows with NULL and/or EMPTY values on the indexed fields will not be indexed.

```
CREATE INDEX nonull_phone ON baggageInfo (contactPhone) WITH NO NULLS
```

- The above query creates an index on the phone number of the passengers. If some passengers do not have a phone number then those fields will not be part of the index.
- 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. 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.

Create an index with unique keys per row

You can create an index with unique keys per row property.

```
CREATE INDEX idx_showid ON
stream_acct(acct_data.contentStreamed[].showId AS INTEGER)
WITH UNIQUE KEYS PER ROW
```

In the above query, an index is created on `showId` and there cannot be duplicate `showId` for a single `contentStreamed` array. This informs the query processor that for any streaming user, the `contentStreamed` array cannot contain two or more shows with the same show id. The restriction is necessary because if duplicate show ids existed, they wouldn't be included in the index. If you insert a row with the same `showId` two or more items in a single `contentStreamed` array, an error is thrown and the insert operation is not successful.

Optimization in the query run time :

When you create an index with unique keys per row, the index would contain fewer entries than the number of elements in the `contentStreamed` array. You could write an efficient query to use this index. The use of such an index by the query would yield fewer results from the FROM clause than if the index was not used.

Examples of creating indexes on functions:

Example 1: Create an index which indexes the rows of the `BaggageInfo` table by their latest modification time:

```
CREATE INDEX idx_modtime ON BaggageInfo(modification_time())
```

This index will be used in a query which has `modification_time` as the filter condition.

```
SELECT * FROM BaggageInfo $u WHERE
modification_time($u) > "2019-08-01T10:45:00"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

This query returns all the rows whose most recent modification time is after 2019-08-01T10:45:00. It uses the `idx_modtime` index defined above. You can verify this by viewing the query plan using the `show query` command.

Example 2: Create an index which indexes the rows of the `BaggageInfo` table on the length of the routing field.

```
CREATE INDEX idx_routlen ON BaggageInfo (length(bagInfo[].routing as string))
```

This index will be used in a query which has `length` as the filter condition.

```
SELECT * from BaggageInfo $bag where length($bag.bagInfo[].routing) > 10
```

This query returns all the rows whose length of the routing field is greater than 10. It uses the `idx_routlen` index defined above. You can verify this by viewing the query plan using the `show query` command.

Example 3: 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",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        "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"
      }
    ]
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Using TableRequest API

You can use TableRequest API to create an index on a NoSQL table.

- Java
- Python
- Go
- Node.js
- C#

Java

The `TableRequest` class is used to create an index on a table. Execution of operations specified by this request is asynchronous. These are potentially long-running operations. `TableResult` is returned from `TableRequest` operations and it encapsulates the state of the table. See Oracle NoSQL Java SDK API Reference for more details on the `TableRequest` class and its methods.

Download the full code ***Indexes.java*** from the examples here.

```
/**
* Create an index acct_episodes in the stream_acct table
*/
private static void crtIndex(NoSQLHandle handle) throws Exception {
   String createIndexDDL = "CREATE INDEX acct_episodes ON " + tableName +

"(acct_data.contentStreamed[].seriesInfo[].episodes[]  AS ANYATOMIC)";
   TableRequest treq = new TableRequest().setStatement(createIndexDDL);
   TableResult tres = handle.tableRequest(treq);
   tres.waitForCompletion(handle, 60000, /* wait 60 sec */
          1000); /* delay ms for poll */
   System.out.println("Index acct_episodes on " + tableName + " is created");
}
```

Python

The `borneo.TableRequest` class is used to create an index on a table. All calls to `borneo.NoSQLHandle.table_request()` are asynchronous so it is necessary to check the result and call `borneo.TableResult.wait_for_completion()` to wait for the operation to complete. See Oracle NoSQL Python SDK API Reference for more details on `table_request` and its methods.

Download the full code ***Indexes.py*** from the examples here.

```
#create an index
def create_index(handle):
   statement = '''CREATE INDEX acct_episodes ON stream_acct
(acct_data.contentStreamed[].seriesInfo[].episodes[]  AS ANYATOMIC)'''
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   request = TableRequest().set_statement(statement)
   table_result = handle.do_table_request(request, 40000, 3000)
   table_result.wait_for_completion(handle, 40000, 3000)
   print('Index acct_episodes on the table stream_acct is created')
```

Go

The `TableRequest` class is used to create an index on a table. Execution of operations specified by `TableRequest` is asynchronous. These are potentially long-running operations. This request is used as the input of a `Client.DoTableRequest()` operation, which returns a `TableResult` that can be used to poll until the table reaches the desired state. See Oracle NoSQL Go SDK API Reference for more details on the various methods of the `TableRequest` class.

Download the full code ***Indexes.go*** from the examples here.

```
//create an index on a table
func createIndex(client *nosqldb.Client, err error, tableName string)(){
   stmt := fmt.Sprintf("CREATE INDEX acct_episodes ON %s "+
        "(acct_data.contentStreamed[].seriesInfo[].episodes[]  AS 
ANYATOMIC)",tableName)
   tableReq := &nosqldb.TableRequest{
        Statement: stmt,
   }
   tableRes, err := client.DoTableRequest(tableReq)
   if err != nil {
      fmt.Printf("cannot initiate CREATE INDEX request: %v\n", err)
      return
   }
   // The create index request is asynchronous, wait for index creation to 
complete.
   _, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)
   if err != nil {
      fmt.Printf("Error finishing CREATE INDEX request: %v\n", err)
      return
   }
   fmt.Println("Created Index acct_episodes on table ", tableName)
   return
}
```

Node.js

You can create an index on a table using the `tableDDL` method. This method is asynchronous and it returns a Promise of `TableResult`. The `TableResult` is a plain JavaScript object that encapsulates the state of the table. For method details, see NoSQLClient class.

Download the full JavaScript code ***Indexes.js*** from the examples here and the full TypeScript code ***Indexes.ts*** from the examples here.

```
//creates an index
async function createIndex(handle) {
   const crtindDDL = `CREATE INDEX acct_episodes ON ${TABLE_NAME}
(acct_data.contentStreamed[].seriesInfo[].episodes[]  AS ANYATOMIC)`;
   let res =  await handle.tableDDL(crtindDDL);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   console.log('Index acct_episodes is created on table:' + TABLE_NAME);
}
```

C#

To create an index on a table use either of the methods `ExecuteTableDDLAsync` or `ExecuteTableDDLWithCompletionAsync`. Both these methods return `Task<TableResult>`. `TableResult` instance contains status of DDL operation such as TableState and table schema. See Oracle NoSQL Dotnet SDK API Reference for more details on these methods.

Download the full code ***Indexes.cs*** from the examples here.

```
// Creates an index on a table
private static async Task createIndex(NoSQLClient client){
   var sql =
      $@"CREATE INDEX acct_episodes ON {TableName}
(acct_data.contentStreamed[].seriesInfo[].episodes[]  AS ANYATOMIC)";
   var tableResult = await client.ExecuteTableDDLAsync(sql);
   // Wait for the operation completion
   await tableResult.WaitForCompletionAsync();
   Console.WriteLine(" Index acct_episodes is created on table Table {0}",
                tableResult.TableName);
}
```

View Index

You can view the indexes in your database.

SHOW INDEXES

The SHOW INDEXES statement provides the list of indexes present in the specified table. If you want the output to be in JSON format, you can specify the optional AS JSON.

Example 1: List indexes on the `BaggageInfo` table.

```
SHOW INDEXES ON baggageInfo

indexes
  jsonindex_routing
  jsonindex_tagnum
  simpleindex_arrival
  nonull_phone
```

Example 2: List indexes on the `BaggageInfo` table in JSON format.

```
SHOW AS JSON INDEXES ON baggageInfo
{"indexes"  :
["jsonindex_routing","jsonindex_tagnum","simpleindex_arrival"]}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

DESCRIBE INDEX

The DESCRIBE INDEX statement defines the specified index on a table. 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 1: Describe the index `multikeyindex3`.

```
DESCRIBE INDEX multikeyindex3 ON stream_acct
+-------------+---------------+-----------+----------
+----------------------------------+-------------+-------------+
table        |  name           | type      | multiKey |
fields                                                | declaredType |
description
+-------------+---------------+-----------+----------
+----------------------------------+-------------+-------------+
stream_acct | multikeyindex3 | SECONDARY | Y        |
acct_data.country                                    | ANY_ATOMIC
|                                 | |               |
            |                 |           |         |
acct_data.contentStreamed[].seriesInfo[].episodes[]| ANY_ATOMIC
|
+-------------+---------------+-----------+----------
+----------------------------------+-------------+-------------+
```

Example 2: Describe the index `idx_showid_year_month` in JSON format.

```
DESCRIBE AS JSON INDEX idx_showid_year_month ON stream_acct
{
  "name" : "idx_showid_year_month",
  "type" : "secondary",
  "fields" : ["acct_data.contentStreamed[].showId",
"substring#acct_data.contentStreamed[].seriesInfo[].episodes[].date@,0,4",
"substring#acct_data.contentStreamed[].seriesInfo[].episodes[].date@,5,2"],
  "types" : ["INTEGER", "STRING", "STRING"],
  "withNoNulls" : false,
  "withUniqueKeysPerRow" : false
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

3 Manage

The articles in this section provide steps on how to manage various database objects.

Namespace Management

A namespace defines a group of tables, within which all of the table names must be uniquely identified. Namespaces permit you to do table privilege management as a group operation.

- Namespace Resolution
- Manage Namespaces
- Namespace scoped privileges
- Granting Authorization Access to Namespaces

Namespace Resolution

You can grant authorization permissions to a namespace to determine who can access both the namespace and the tables within it.

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.
  - 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.

Manage Namespaces

SHOW NAMESPACES

The SHOW NAMESPACES statement provides the list of namespaces in the system. You can specify **AS JSON** if you want the output to be in JSON format.

Example 1: The following statement lists the namespaces present in the system.

```
SHOW NAMESPACES
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
namespaces
  sysdefault
```

Example 2: The following statement lists the namespaces present in the system in JSON format.

```
SHOW AS JSON NAMESPACES
```

Output:

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

DROP NAMESPACE

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

IF EXISTS 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 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.

The following statement removes the namespace named **ns1**.

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

Using APIs to drop namespaces:

- Java
- Python
- Go
- Node.js
- C#

Java

You can drop a namespace using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation. These operations are asynchronous and completion needs to be checked.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***Namespaces.java*** from the examples here.

```
private static void dropNS(NoSQLHandle handle) throws Exception {
    String dropNSDDL = "DROP NAMESPACE " + nsName;
    SystemRequest sysreq = new SystemRequest();
    sysreq.setStatement(dropNSDDL.toCharArray());
    SystemResult sysres = handle.systemRequest(sysreq);
    sysres.waitForCompletion(handle, 60000,1000);
    System.out.println("Namespace " + nsName + " is dropped");
}
```

Python

You can drop a namespace using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation.

Download the full code ***Namespaces.py*** from the examples here.

```
def drop_ns(handle):
   statement = '''DROP NAMESPACE ns1'''
   sysreq = SystemRequest().set_statement(statement)
   sys_result = handle.system_request(sysreq)
   sys_result.wait_for_completion(handle, 40000, 3000)
   print('Namespace: ns1 is dropped')
```

Go

You can drop a namespace using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operations. These are potentially long-running operations and completion of the operation needs to be checked.

Download the full code ***Namespaces.go*** from the examples here.

```
func dropNS(client *nosqldb.Client, err error)(){
    stmt := fmt.Sprintf("DROP NAMESPACE ns1")
    sysReq := &nosqldb.SystemRequest{
        Statement: stmt,
    }
    sysRes, err := client.DoSystemRequest(sysReq)
    _, err = sysRes.WaitForCompletion(client, 60*time.Second, time.Second)
    if err != nil {
        fmt.Printf("Error finishing CREATE NAMESPACE request: %v\n", err)
        return
    }
    fmt.Println("Dropped Namespace ns1 ")
    return
}
```

Node.js

You can create namespace using `adminDDL` method. The `adminDDL` method is used to perform any table-independent administrative operation.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full JavaScript code ***Namespaces.js*** from here and the full TypeScript code ***Namespaces.ts*** from here.

```
async function dropNS(handle) {
   const dropNS = `DROP NAMESPACE ns1`;
   let res = await handle.adminDDL(dropNS);
   console.log('Namespace dropped: ns1' );
}
```

C#

The `ExecuteAdminSync` method is used to perform any table-independent administrative operations.

Download the full code ***Namespaces.cs*** from the examples here.

```
private static async Task dropNS(NoSQLClient client){
   var sql = $@"DROP NAMESPACE ns1";
   var adminResult = await client.ExecuteAdminAsync(sql);
   // Wait for the operation completion
   await adminResult.WaitForCompletionAsync();
   Console.WriteLine("Dropped namespace ns1");
}
```

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>;`

Copyright © 2022, 2025, Oracle and/or its affiliates.

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

<table>
<tr><th>Privilege</th><th>Description</th></tr>
<tr><td>CREATE_TABLE_IN_NAMESPACE
DROP_TABLE_IN_NAMESPACE
EVOLVE_TABLE_IN_NAMESPACE</td><td>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.

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;</td></tr>
<tr><td>CREATE_INDEX_IN_NAMESPACE
DROP_INDEX_IN_NAMESPACE</td><td>Grant permission to a user or to a role to create or drop an index in a specific namespace.

GRANT CREATE_INDEX_IN_NAMESPACE ON NAMESPACE namespace_name TO <User|Role>;
GRANT DROP_INDEX_IN_NAMESPACE ON NAMESPACE namespace_name TO <User|Role>;</td></tr>
<tr><td>READ_IN_NAMESPACE
INSERT_IN_NAMESPACE
DELETE_IN_NAMESPACE</td><td>Grant permission to a role to read, insert, or delete items in a specific namespace.

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>;</td></tr>
<tr><td>MODIFY_IN_NAMESPACE</td><td>Helper label for granting or revoking permissions to all DDL privileges for a specific namespace to a user or role.

GRANT MODIFY_IN_NAMESPACE ON NAMESPACE namespace_name TO <User|Role>;
REVOKE MODIFY_IN_NAMESPACE ON NAMESPACE namespace_name TO <User|Role>;</td></tr>
</table>

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]
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

```
Example:

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.

```
Example:
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
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.

Managing Tables, Indexes & Regions

You will learn different ways to alter an existing table. You will also learn how to manage indexes and regions.

- Alter Table
- Drop Table
- Drop Index
- Manage regions

Alter Table

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

Copyright © 2022, 2025, Oracle and/or its affiliates.

- Modify the Time-To-Live value of the table
- Add a region
- Remove a region

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.

- Using SQL command to alter table
- Using TableRequest API to alter table

Using SQL command to alter table

You can use ALTER TABLE command to change the definition of the table.

Create a sample table :

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

Add/remove schema fields

Example : Add schema field to the table schema.

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

Explanation: 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 maybe 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.

Example : Remove schema fields in the table schema.

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

Explanation: 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
```

Alter TTL value

Copyright © 2022, 2025, Oracle and/or its affiliates.

Example : 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 no longer available. 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 command to change this value for any table. You can specify the TTL with a number, followed by either `HOURS` or `DAYS`.

```
ALTER TABLE stream_acct USING TTL 5 days
```

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 placed in rows created subsequent to the alter 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.

Add/remove a region

Example : 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.

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

```
ALTER TABLE myTable ADD REGIONS FRA;
```

Explanation: Here, `myTable` is an MR table and `FRA` is an existing region.

Example : 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 NoSQL Database setup, you must run the following steps from all the other participating regions.

```
ALTER TABLE myTable DROP REGIONS FRA
```

Here, `myTable` is a MR Table and `FRA` is the region to be dropped. You can supply a comma_separated_list_of_regions if you want to drop more than one region.

Enable before-images

Example : Enable before-images for table write operations

The before-image of any write is the table row before it gets updated or deleted by a DML operation You can use the ALTER TABLE statement with the `ENABLE BEFORE IMAGE` clause to

Copyright © 2022, 2025, Oracle and/or its affiliates.

enable the generation of before-images for any write operations on the table as shown in the example below:

```
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.

Note

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.

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.

Using TableRequest API to alter table

You can use TableRequest API to change the definition of a NoSQL table.

- Java
- Python
- Go
- Node.js
- C#

Java

The `TableRequest` class is used to modify tables. Execution of operations specified by this request is asynchronous. These are potentially long-running operations. `TableResult` is returned from `TableRequest` operations and it encapsulates the state of the table. See Oracle NoSQL Java SDK API Reference for more details on the `TableRequest` class and its methods.

Download the full code ***AlterTable.java*** from the examples here.

```
/**
* Alter the table stream_acct and add a column
*/
private static void alterTab(NoSQLHandle handle) throws Exception {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    String alterTableDDL = "ALTER TABLE " + tableName +"(ADD acctname STRING)";
    TableRequest treq = new TableRequest().setStatement(alterTableDDL);
    System.out.println("Altering table " + tableName);
    TableResult tres = handle.tableRequest(treq);
    tres.waitForCompletion(handle, 60000, /* wait 60 sec */
    1000); /* delay ms for poll */
    System.out.println("Table " + tableName + " is altered");
}
```

Python

The `borneo.TableRequest` class is used to modify tables. All calls to `borneo.NoSQLHandle.table_request()` are asynchronous so it is necessary to check the result and call `borneo.TableResult.wait_for_completion()` to wait for the operation to complete. See Oracle NoSQL Python SDK API Reference for more details on `table_request` and its methods.

Download the full code ***AlterTable.py*** from the examples here.

```
def alter_table(handle):
   statement = '''ALTER TABLE  stream_acct(ADD acctname STRING)'''
   request = TableRequest().set_statement(statement)
   table_result = handle.do_table_request(request, 40000, 3000)
   table_result.wait_for_completion(handle, 40000, 3000)
   print('Table stream_acct is altered')
```

Go

The `TableRequest` class is used to modify tables. Execution of operations specified by `TableRequest` is asynchronous. These are potentially long-running operations. This request is used as the input of a `Client.DoTableRequest()` operation, which returns a `TableResult` that can be used to poll until the table reaches the desired state. See Oracle NoSQL Go SDK API Reference for more details on the various methods of the `TableRequest` class.

Download the full code ***AlterTable.go*** from the examples here.

```
//alter an existing table and add a column
func alterTable(client *nosqldb.Client, err error, tableName string)(){
   stmt := fmt.Sprintf("ALTER TABLE %s (ADD acctName STRING)",tableName)
   tableReq := &nosqldb.TableRequest{
        Statement: stmt,
   }
   tableRes, err := client.DoTableRequest(tableReq)
   if err != nil {
      fmt.Printf("cannot initiate ALTER TABLE request: %v\n", err)
      return
   }
   // The alter table request is asynchronous, wait for table alteration to
complete.
   _, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)
   if err != nil {
      fmt.Printf("Error finishing ALTER TABLE request: %v\n", err)
      return
   }
   fmt.Println("Altered table ", tableName)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   return
}
```

Node.js

You can use execute the `tableDDL` method to modify a table. This method is asynchronous and it returns a Promise of `TableResult`. The `TableResult` is a plain JavaScript object that encapsulates the state of the table after the DDL operation. For method details, see NoSQLClient class.

Download the full JavaScript code ***AlterTable.js*** from the examples here and the full TypeScript code ***AlterTable.ts*** from the examples here.

```
//alter a table and add a column
async function alterTable(handle) {
   const alterDDL = `ALTER TABLE ${TABLE_NAME} (ADD acctname STRING)`;
   let res =  await handle.tableDDL(alterDDL);
   console.log('Table altered: ' + TABLE_NAME);
}
```

C#

You can use either of the two methods `ExecuteTableDDLAsync` and `ExecuteTableDDLWithCompletionAsync` to modify a table. Both the methods return `Task<TableResult>`. `TableResult` instance encapsulates the state of the table after the DDL operation. See Oracle NoSQL Dotnet SDK API Reference for more details on these methods.

Download the full code ***AlterTable.cs*** from the examples here.

```
private static async Task alterTable(NoSQLClient client){
   var sql = $@"ALTER TABLE {TableName}(ADD acctname STRING)";
   var tableResult = await client.ExecuteTableDDLAsync(sql);
   // Wait for the operation completion
   await tableResult.WaitForCompletionAsync();
   Console.WriteLine("  Table {0} is altered", tableResult.TableName);
}
```

Drop Table

- Using SQL command to drop table
- Using TableRequest API to drop table

Using SQL command to drop table

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

By default, if the named table does not exist then this statement fails. You don't get an error if the optional IF EXISTS clause is specified and the table does not exist.

```
DROP TABLE demo_acct
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Note

To drop a MR Table, first drop all of its child tables. Otherwise, the DROP statement results in an error.

Using TableRequest API to drop table

You can use TableRequest API to drop a NoSQL table.

- Java
- Python
- Go
- Node.js
- C#

Java

Execution of operations specified by`TableRequest` is asynchronous. These are potentially long-running operations. `TableResult` is returned from `TableRequest` operations and it encapsulates the state of the table. See Oracle NoSQL Java SDK API Reference for more details on the `TableRequest` class and its methods.

Download the full code ***AlterTable.java*** from the examples here.

```
/*Drop the table*/
private static void dropTab(NoSQLHandle handle) throws Exception {
   String dropTableDDL = "DROP TABLE " + tableName;
   TableRequest treq = new TableRequest().setStatement(dropTableDDL);
   TableResult tres = handle.tableRequest(treq);
   tres.waitForCompletion(handle, 60000, /* wait 60 sec */
       1000); /* delay ms for poll */
   System.out.println("Table " + tableName + " is dropped");
}
```

Python

You can use the `borneo.TableRequest` class to drop a table. All calls to `borneo.NoSQLHandle.table_request()` are asynchronous so it is necessary to check the result and call `borneo.TableResult.wait_for_completion()` to wait for the operation to complete. See Oracle NoSQL Python SDK API Reference for more details on `table_request` and its methods.

Download the full code ***AlterTable.py*** from the examples here.

```
def drop_table(handle):
   statement = '''DROP TABLE stream_acct'''
   request = TableRequest().set_statement(statement)
   table_result = handle.do_table_request(request, 40000, 3000)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    table_result.wait_for_completion(handle, 40000, 3000)
    print('Dropped table: stream_acct')
```

Go

You can use the `TableRequest` class to drop a table. Execution of operations specified by `TableRequest` is asynchronous. These are potentially long-running operations. This request is used as the input of a `Client.DoTableRequest()` operation, which returns a `TableResult` that can be used to poll until the table reaches the desired state. See Oracle NoSQL Go SDK API Reference for more details on the various methods of the `TableRequest` class.

Download the full code ***AlterTable.go*** from the examples here.

```
//drop an existing table
func dropTable(client *nosqldb.Client, err error, tableName string)(){
   stmt := fmt.Sprintf("DROP TABLE %s",tableName)
   tableReq := &nosqldb.TableRequest{
        Statement: stmt,
   }
   tableRes, err := client.DoTableRequest(tableReq)
   return
}
```

Node.js

You can use the `tableDDL` method to drop a table. This method is asynchronous and it returns a Promise of `TableResult`. The `TableResult` is a plain JavaScript object that encapsulates the state of the table after the DDL operation. For method details, see NoSQLClient class.

Download the full JavaScript code ***AlterTable.js*** from the examples here and the full TypeScript code ***AlterTable.ts*** from the examples here.

```
//drop a table
async function dropTable(handle) {
   const dropDDL = `DROP TABLE ${TABLE_NAME}`;
   let res =  await handle.tableDDL(dropDDL);
   console.log('Table dropped: ' + TABLE_NAME);
}
```

C#

You can use either of the methods `ExecuteTableDDLAsync` or `ExecuteTableDDLWithCompletionAsync` to drop a table. Both the methods return `Task<TableResult>`. `TableResult` instance encapsulates the state of the table after the DDL operation. See Oracle NoSQL Dotnet SDK API Reference for more details on these methods.

Download the full code ***AlterTable.cs*** from the examples here.

```
private static async Task dropTable(NoSQLClient client){
   var sql = $@"DROP TABLE {TableName}";
   var tableResult = await client.ExecuteTableDDLAsync(sql);
   // Wait for the operation completion
   await tableResult.WaitForCompletionAsync();
   Console.WriteLine("  Table {0} is dropped", tableResult.TableName);
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Drop Index

You can drop an index from your database when you no longer need it.

- Using SQL command to drop index
- Using TableRequest API to drop index

Using SQL command to drop index

The DROP INDEX removes the specified index from the database.

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: Drop the index `multikeyindex1`.

```
DROP INDEX multikeyindex1 ON stream_acct
```

Using TableRequest API to drop index

You can use TableRequest API to drop an index of a NoSQL table.

- Java
- Python
- Go
- Node.js
- C#

Java

Execution of operations specified by `TableRequest` class is asynchronous. These are potentially long-running operations. `TableResult` is returned from `TableRequest` operations and it encapsulates the state of the table. See Oracle NoSQL Java SDK API Reference for more details on the TableRequest class and its methods.

Download the full code ***Indexes.java*** from the examples here.

```
/* Drop the index acct_episodes*/
private static void dropIndex(NoSQLHandle handle) throws Exception {
   String dropIndexDDL = "DROP INDEX acct_episodes ON " + tableName;
   TableRequest treq = new TableRequest().setStatement(dropIndexDDL);
   TableResult tres = handle.tableRequest(treq);
   tres.waitForCompletion(handle, 60000, /* wait 60 sec */
            1000); /* delay ms for poll */
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   System.out.println("Index acct_episodes on " + tableName + " is dropped");
}
```

Python

You can use the `borneo.TableRequest` class to drop a table index. All calls to `borneo.NoSQLHandle.table_request()` are asynchronous so it is necessary to check the result and call `borneo.TableResult.wait_for_completion()` to wait for the operation to complete. See Oracle NoSQL Python SDK API Reference for more details on `table_request` and its methods.

Download the full code ***Indexes.py*** from the examples here.

```
#drop the index
def drop_index(handle):
   statement = '''DROP INDEX acct_episodes ON stream_acct'''
   request = TableRequest().set_statement(statement)
   table_result = handle.do_table_request(request, 40000, 3000)
   table_result.wait_for_completion(handle, 40000, 3000)
   print('Index acct_episodes on the table stream_acct is dropped')
```

Go

You can use the `TableRequest` class to drop a table index. Execution of operations specified by `TableRequest` is asynchronous. These are potentially long-running operations. This request is used as the input of a `Client.DoTableRequest()` operation, which returns a `TableResult` that can be used to poll until the table reaches the desired state. See Oracle NoSQL Go SDK API Reference for more details on the various methods of the `TableRequest` class.

Download the full code ***Indexes.go*** from the examples here.

```
//drops an index from a table
func dropIndex(client *nosqldb.Client, err error, tableName string)(){
   stmt := fmt.Sprintf("DROP INDEX acct_episodes ON %s",tableName)
   tableReq := &nosqldb.TableRequest{
        Statement: stmt,
   }
   tableRes, err := client.DoTableRequest(tableReq)
   if err != nil {
      fmt.Printf("cannot initiate DROP INDEX request: %v\n", err)
      return
   }
   // The drop index request is asynchronous, wait for drop index to complete.
   _, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)
   if err != nil {
      fmt.Printf("Error finishing DROP INDEX request: %v\n", err)
      return
   }
   fmt.Println("Dropped index acct_episodes on table ", tableName)
   return
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Node.js

You can use the `tableDDL` method to drop a table index. This method is asynchronous and it returns a Promise of `TableResult`. The `TableResult` is a plain JavaScript object that contains the status of the DDL operation such as its TableState, name, schema, and its TableLimits. For method details, see NoSQLClient class.

Download the full JavaScript code ***Indexes.js*** from the examples here and the full TypeScript code ***Indexes.ts*** from the examples here.

```
//drops an index
async function dropIndex(handle) {
   const dropindDDL = `DROP INDEX acct_episodes ON ${TABLE_NAME}`;
   let res =  await handle.tableDDL(dropindDDL);
   console.log('Index acct_episodes is dropped from table: ' + TABLE_NAME);
}
```

C#

You can use one of the methods `ExecuteTableDDLAsync` or `ExecuteTableDDLWithCompletionAsync` to drop a table index. Both the methods return `Task<TableResult>`. `TableResult` instance contains status of DDL operation such as TableState and table schema. See Oracle NoSQL Dotnet SDK API Reference for more details on these methods.

Download the full code ***Indexes.cs*** from the examples here.

```
private static async Task dropIndex(NoSQLClient client){
   var sql = $@"DROP INDEX acct_episodes on {TableName}";
   var tableResult = await client.ExecuteTableDDLAsync(sql);
   // Wait for the operation completion
   await tableResult.WaitForCompletionAsync();
   Console.WriteLine(" Index acct_episodes is dropped from table Table {0}",
                tableResult.TableName);
}
```

Manage regions

The `show regions` statement provides the list of regions present in the Multi-Region Oracle NoSQL Database. You need to specify "AS JSON" if you want the output to be in JSON format.

Example 1: The following statement lists all the existing regions.

```
SHOW REGIONS
```

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

```
SHOW AS JSON REGIONS
```

In a Multi-Region Oracle NoSQL Database environment, the drop region statement removes the specified remote region from the local region. See Set up Multi-Region Environment for more details on the local regions and remote regions in a Multi-Region setup.

Copyright © 2022, 2025, Oracle and/or its affiliates.

> **Note**
>
> This region must be different from the local region where the command is executed.

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

```
DROP REGION my_region1
```

Using APIs to drop regions:

- Java
- Python
- Go
- Node.js
- C#

Java

You can drop a region using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operation. You can use the `TableRequest` class to drop tables.

Download the full code ***Regions.java*** from the examples here.

```
/* Drop a table from a region*/
private static void dropTabInRegion(NoSQLHandle handle) throws Exception {
   String dropTableDDL = "DROP TABLE " + tableName;
   TableRequest treq = new TableRequest().setStatement(dropTableDDL);
   TableResult tres = handle.tableRequest(treq);
   tres.waitForCompletion(handle, 60000, /* wait 60 sec */
   1000); /* delay ms for poll */
   System.out.println("Table " + tableName + " is dropped");
}
   /* Drop a region*/
   private static void dropRegion(NoSQLHandle handle, String regName) throws
Exception {
      String dropNSDDL = "DROP REGION " + regName;
      SystemRequest sysreq = new SystemRequest();
      sysreq.setStatement(dropNSDDL.toCharArray());
      SystemResult sysres = handle.systemRequest(sysreq);
      sysres.waitForCompletion(handle, 60000,1000);
      System.out.println("Region " + regName + " is dropped");
   }
```

Python

You can drop a region using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operations. You can drop a table using the `borneo.TableRequest` class.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***Regions.py*** from the examples here.

```
#Drop the table from a region
def drop_tab_region(handle):
   statement = '''DROP TABLE stream_acct'''
   request = TableRequest().set_statement(statement)
   table_result = handle.do_table_request(request, 40000, 3000)
   table_result.wait_for_completion(handle, 40000, 3000)
   print('Dropped table: stream_acct')

#Drop the region
def drop_region(handle):
   statement = '''DROP REGION LON'''
   sysreq = SystemRequest().set_statement(statement)
   sys_result = handle.system_request(sysreq)
   sys_result.wait_for_completion(handle, 40000, 3000)
   print('Region LON is dropped')
```

Go

You can drop a region using `SystemRequest` class. The `SystemRequest` class is used to perform any table-independent administrative operations. You can drop a table using `TableRequest` class.

Download the full code ***Regions.go*** from the examples here.

```
//drops a table from a region
func drpTabInRegion(client *nosqldb.Client, err error, tableName string)(){
   stmt := fmt.Sprintf("DROP TABLE %s",tableName)
   tableReq := &nosqldb.TableRequest{
        Statement: stmt,
   }
   tableRes, err := client.DoTableRequest(tableReq)
   if err != nil {
      fmt.Printf("cannot initiate DROP TABLE request: %v\n", err)
      return
   }
   _, err = tableRes.WaitForCompletion(client, 60*time.Second, time.Second)
   if err != nil {
      fmt.Printf("Error finishing DROP TABLE request: %v\n", err)
      return
   }
   fmt.Println("Dropped table ", tableName)
   return
}

//drop a region
func dropRegion(client *nosqldb.Client, err error)(){
   stmt := fmt.Sprintf("DROP REGION LON")
   sysReq := &nosqldb.SystemRequest{
        Statement: stmt,
   }
   sysRes, err := client.DoSystemRequest(sysReq)
   _, err = sysRes.WaitForCompletion(client, 60*time.Second, time.Second)
   if err != nil {
      fmt.Printf("Error finishing DROP REGION request: %v\n", err)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        return
    }
    fmt.Println("Dropped REGION LON ")
    return
}
```

Node.js

You can drop a region using `adminDDL` method. The `adminDDL` module is used to perform an administrative operation on the system. You can drop a table using the `tableDDL` method.

Download the full JavaScript code ***Regions.js*** from the examples here and the full TypeScript code ***Regions.ts*** from the examples here.

```
//drop a table from a region
async function dropTabInRegion(handle) {
   const dropDDL = `DROP TABLE ${TABLE_NAME}`;
   let res =  await handle.tableDDL(dropDDL);
   console.log('Table dropped: ' + TABLE_NAME);
}

//drop a region
async function dropRegion(handle) {
   const dropReg = `DROP REGION LON`;
   let res = await handle.adminDDL(dropReg);
   console.log('Region dropped: LON' );
}
```

C#

You can use `ExecuteAdminSync` method to drop a region. The `ExecuteAdminSync` method is used to perform an administrative operation on the system. You can drop a table using either `ExecuteTableDDLAsync` or `ExecuteTableDDLWithCompletionAsync`.

Download the full code ***Regions.cs*** from the examples here.

```
private static async Task dropTabInRegion(NoSQLClient client){
   var sql = $@"DROP TABLE {TableName}";
   var tableResult = await client.ExecuteTableDDLAsync(sql);
   // Wait for the operation completion
   await tableResult.WaitForCompletionAsync();
   Console.WriteLine("  Table {0} is dropped", tableResult.TableName);
}

private static async Task dropRegion(NoSQLClient client){
   var sql = $@"DROP REGION LON";
   var adminResult = await client.ExecuteAdminAsync(sql);
   // Wait for the operation completion
   await adminResult.WaitForCompletionAsync();
   Console.WriteLine("  Dropped region LON");
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

4

Develop

The articles in this section provide steps on how to use SQL and write queries. It covers information about different complex data types. It also covers how to use indexes for query optimization.

Inserting, Modifying, and Deleting Data

You can perform various data manipulation operations in your table. You can add data, modify an existing data and remove data.

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Insert data
- Upsert Data
- Update Data
- Modify JSON data
- Delete Data

Insert data

- Using SQL command to insert data
- Using Put API to insert data
- Using MultiWrite API to insert data

Using SQL command to insert data

The INSERT statement is used to construct a new row and add it to a specified table.

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). These expressions/DEFAULTs compute the value for their associated column in the new row. 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 1: Inserting a row into `BaggageInfo` table providing all column values:

```
INSERT INTO BaggageInfo VALUES(
1762392196147,
"Birgit Naquin",
"M",
"165-742-5715",
"QD1L0T",
[ {
    "id" : "7903989918469",
    "tagNum" : "17657806240229",
    "routing" : "JFK/MAD",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "MAD",
    "flightLegs" : [ {
      "flightNo" : "BM495",
      "flightDate" : "2019-03-07T07:00:00Z",
      "fltRouteSrc" : "JFK",
      "fltRouteDest" : "MAD",
      "estimatedArrival" : "2019-03-07T14:00:00Z",
      "actions" : [ {
        "actionAt" : "MAD",
        "actionCode" : "Offload to Carousel at MAD",
        "actionTime" : "2019-03-07T13:54:00Z"
      }, {
        "actionAt" : "JFK",
        "actionCode" : "ONLOAD to MAD",
        "actionTime" : "2019-03-07T07:00:00Z"
      }, {
        "actionAt" : "JFK",
        "actionCode" : "BagTag Scan at JFK",
        "actionTime" : "2019-03-07T06:53:00Z"
      }, {
        "actionAt" : "JFK",
        "actionCode" : "Checkin at JFK",
        "actionTime" : "2019-03-07T05:03:00Z"
      } ]
    } ],
    "lastSeenTimeGmt" : "2019-03-07T13:51:00Z",
    "bagArrivalDate" : "2019-03-07T13:51:00Z"
  } ]
)
```

Example 2: Skipping some data while doing an INSERT statement by specifying the DEFAULT clause.

You can skip the data of some columns by specifying "DEFAULT".

```
INSERT INTO BaggageInfo VALUES(
1762397286805,
"Bonnie Williams",
DEFAULT,
DEFAULT,
"CZ1O5I",
[ {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
      "id" : "79039899129693",
      "tagNum" : "17657806216554",
      "routing" : "SFO/ORD/FRA",
      "lastActionCode" : "OFFLOAD",
      "lastActionDesc" : "OFFLOAD",
      "lastSeenStation" : "FRA",
      "flightLegs" : [ {
        "flightNo" : "BM572",
        "flightDate" : "2019-03-02T05:00:00Z",
        "fltRouteSrc" : "SFO",
        "fltRouteDest" : "ORD",
        "estimatedArrival" : "2019-03-02T09:00:00Z",
        "actions" : [ {
          "actionAt" : "SFO",
          "actionCode" : "ONLOAD to ORD",
          "actionTime" : "2019-03-02T05:24:00Z"
        }, {
          "actionAt" : "SFO",
          "actionCode" : "BagTag Scan at SFO",
          "actionTime" : "2019-03-02T04:52:00Z"
        }, {
          "actionAt" : "SFO",
          "actionCode" : "Checkin at SFO",
          "actionTime" : "2019-03-02T03:28:00Z"
        } ]
      }, {
        "flightNo" : "BM582",
        "flightDate" : "2019-03-02T05:24:00Z",
        "fltRouteSrc" : "ORD",
        "fltRouteDest" : "FRA",
        "estimatedArrival" : "2019-03-02T13:24:00Z",
        "actions" : [ {
          "actionAt" : "FRA",
          "actionCode" : "Offload to Carousel at FRA",
          "actionTime" : "2019-03-02T13:20:00Z"
        }, {
          "actionAt" : "ORD",
          "actionCode" : "ONLOAD to FRA",
          "actionTime" : "2019-03-02T12:54:00Z"
        }, {
          "actionAt" : "ORD",
          "actionCode" : "OFFLOAD from ORD",
          "actionTime" : "2019-03-02T12:30:00Z"
        } ]
      } ],
      "lastSeenTimeGmt" : "2019-03-02T13:18:00Z",
      "bagArrivalDate" : "2019-03-02T13:18:00Z"
    } ]
)
```

Example 3: Specifying column names and skipping columns in the insert statement.

Copyright © 2022, 2025, Oracle and/or its affiliates.

If you have data only for some columns of a table, you can specify the name of the columns in the INSERT clause and then specify the corresponding values in the "VALUES" clause.

```
INSERT INTO BaggageInfo(ticketNo, fullName,confNo,bagInfo) VALUES(
1762355349471,
"Bryant Weber",
"LI7N1W",
[ {
    "id" : "79039899149056",
    "tagNum" : "17657806234185",
    "routing" : "MEL/LAX/MIA",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "MIA",
    "flightLegs" : [ {
      "flightNo" : "BM114",
      "flightDate" : "2019-03-01T12:00:00Z",
      "fltRouteSrc" : "MEL",
      "fltRouteDest" : "LAX",
      "estimatedArrival" : "2019-03-02T02:00:00Z",
      "actions" : [ {
        "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"
      } ]
    }, {
      "flightNo" : "BM866",
      "flightDate" : "2019-03-01T12:20:00Z",
      "fltRouteSrc" : "LAX",
      "fltRouteDest" : "MIA",
      "estimatedArrival" : "2019-03-02T16:21:00Z",
      "actions" : [ {
        "actionAt" : "MIA",
        "actionCode" : "Offload to Carousel at MIA",
        "actionTime" : "2019-03-02T16:18:00Z"
      }, {
        "actionAt" : "LAX",
        "actionCode" : "ONLOAD to MIA",
        "actionTime" : "2019-03-02T16:12:00Z"
      }, {
        "actionAt" : "LAX",
        "actionCode" : "OFFLOAD from LAX",
        "actionTime" : "2019-03-02T16:02:00Z"
      } ]
    } ],
    "lastSeenTimeGmt" : "2019-03-02T16:09:00Z",
    "bagArrivalDate" : "2019-03-02T16:09:00Z"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
  } ]
)
```

Example 4: Inserting a row into `stream_acct` table providing all column values:

```
INSERT INTO stream_acct VALUES(
1,
"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"
         }]
      },
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        {
           "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"
           }]
        }]
     }]
});
```

Example 5: Insert data into the JSON collection table created for a shopping application.

```
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}])
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

In the above example, you insert the shopper's data by supplying the `contactPhone` as the primary key followed by other details of the shoppers. The shopper's details are stored as a single document. Notice that in JSON collection tables, you do not supply a column name for the document itself and you 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 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 the items wishlisted by the shopper.

Using Put API to insert data

Add rows to your table. When you store data in table rows, your application can easily retrieve, add to, or delete information from a table.

You can use the `PutRequest` class / `put` method to perform unconditional and conditional puts to:

- Overwrite any existing row. Overwrite is the default functionality.
- Succeed only if the row does not exist. Use the `IfAbsent` method in this case.
- Succeed only if the row exists. Use the `IfPresent` method in this case.
- Succeed only if the row exists and the version matches a specific version. Use `IfVersion` method for this case and the `setMatchVersion` method to specify the version to match.

- Java
- Python
- Go
- Node.js
- C#

Java

The `PutRequest` class provides the `setValueFromJson` method which takes a JSON string and uses that to populate a row to insert into the table. The JSON string should specify field names that correspond to the table field names.

Download the full code ***AddData.java*** from the examples here.

```
private static void writeRows(NoSQLHandle handle, MapValue value)
                                                        throws Exception {
   PutRequest putRequest =
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        new PutRequest().setValue(value).setTableName(tableName);
      PutResult putResult = handle.put(putRequest);

      if (putResult.getVersion() != null) {
         System.out.println("Added a row to the stream_acct table");
      } else {
         System.out.println("Put failed");
      }
   }
```

Inserting data into a JSON collection table: You can insert the top-level fields of the document in the JSON collection table using a sequence of `PutRequest` operations. For nested-level JSON fields, you can supply the JSON string in the `putFromJson` operation. You can also use the `createFromJson` method which takes the fields as a JSON string and uses that to populate a row in the table.

```
/*
 * Construct a row for JSON collection table with the following data:
 * {
 *   "id": 1,
 *   "name": "John Doe",
 *   "age": 25,
 *   "college" : {"name" : "Presidency", "branch" : "Biotechnology"}
 *   }
 * }
 */

String tableName = "usersJSON";

MapValue value = new MapValue().put("id", 1)

            .put("name", "John Doe")
            .put("age", 25)
            .putFromJson("college",
                 "{\"name\" : \"Presidency\"," +
                 "\"branch\" : \"Biotechnology\"" +
                 " }", null);

            PutRequest putRequest = new PutRequest()
                .setValue(value)
                .setTableName(tableName);

            PutResult putRes = handle.put(putRequest);
            System.out.println("Put row: " + value + " result=" + putRes);
```

> **Note**
>
> Oracle NoSQL Database also supports user-defined row metadata. You can annotate actual row data with additional information in a JSON string using the `setRowMetadata()` method. To understand how it works, see Using row metadata in Write Operations.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Python

The `borneo.PutRequest` class represents input to the `borneo.NoSQLHandle.put()` method which is used to insert single rows.

You can also add JSON data to your table. In the case of a fixed-schema table the JSON is converted to the target schema. JSON data can be directly inserted into a column of type JSON. The use of the JSON data type allows you to create table data without a fixed schema, allowing more flexible use of the data.

Download the full code ***AddData.py*** from the examples here.

```
def insert_record(handle,table_name,acct_data):
  request = PutRequest().set_table_name(table_name)
                        .set_value_from_json(acct_data)

  handle.put(request)
  print('Added a row to the stream_acct table')
```

Inserting data into a JSON collection table: You can add a row directly into the JSON collection table as a JSON string.

```
from borneo import PutRequest
request =
PutRequest().set_table_name('usersJSON').request.set_value_from_json('{"id":
1, "name": "John Doe", "age": 25, "college" : {"name" : "Presidency",
"branch" : "Biotechnology"}}')
handle.put(request)
```

Go

The `nosqldb.PutRequest` represents an input to the `nosqldb.Put()` function and is used to insert single rows.

The data value provided for a row (in `PutRequest`) is a `*types.MapValue`. The `key` portion of each entry in the `MapValue` must match the column name of target table, and the `value` portion must be a valid value for the column. JSON data can also be directly inserted into a column of type `JSON`. The use of the JSON data type allows you to create table data without a fixed schema, allowing more flexible use of the data.

Download the full code ***AddData.go*** from the examples here.

```
func insertData(client *nosqldb.Client, err error,
                tableName string,value1 *types.MapValue )(){
  putReq := &nosqldb.PutRequest{
        TableName: tableName,
        Value: value1,
  }
  putRes, err := client.Put(putReq)
  if err != nil {
     fmt.Printf("failed to put single row: %v\n", err)
     return
  }
  fmt.Printf("Added a row to the stream_acct table\n")
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Inserting data into a JSON collection table: You can create a map value from JSON data and add the row to the JSON collection table.

```
value, err:=types.NewMapValueFromJSON(`{"id": 1, "name": "John Doe", "age":
25, "college" : {"name" : "Presidency", "branch" : "Biotechnology"}}`)
iferr!=nil {
    return
}
req:=&nosqldb.PutRequest{
    TableName: "usersJSON",
    Value: value,
}
res, err:=client.Put(req)
```

Node.js

You use the `put` method to insert a single row into the table. For method details, see NoSQLClient class.

JavaScript: Download the full code ***AddData.js*** from the examples here.

```
/* Adding 3 records in acct_stream table */
let putResult = await handle.put(TABLE_NAME, JSON.parse(acct1));
let putResult1 = await handle.put(TABLE_NAME, JSON.parse(acct2));
let putResult2 = await handle.put(TABLE_NAME, JSON.parse(acct3));

console.log("Added rows to the stream_acct table");
```

TypeScript: Download the full code ***AddData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}
/* Adding 3 records in acct_stream table */
let putResult = await handle.put<StreamInt>(TABLE_NAME, JSON.parse(acct1));
let putResult1 = await handle.put<StreamInt>(TABLE_NAME, JSON.parse(acct2));
let putResult2 = await handle.put<StreamInt>(TABLE_NAME, JSON.parse(acct3));

console.log("Added rows to the stream_acct table");
```

Inserting data into a JSON collection table: You can add a row into the JSON collection table by supplying a plain JavaScript object with supported JSON types.

```
import { NoSQLClient, ServiceType } from 'oracle-nosqldb';
const client = new NoSQLClient('config.json');
const TABLE_NAME = 'usersJSON';
const record = {id : 1,
    name : 'John Doe',
    age : 25,
    college : {
        name : 'Presidency',
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        branch : 'Biotechnology'
        }
}


async function writeARecord(client, record) {
    await client.put(TABLE_NAME, record);
}
```

C#

The method `PutAsync` and related methods `PutIfAbsentAsync` , `PutIfPresentAsync` and `PutIfVersionAsync` are used to insert a single row into the table or update a single row.

Each of the `Put` methods above returns `Task<PutResult<RecordValue>>`. `PutResult` instance contains info about a completed `Put` operation, such as success status (conditional put operations may fail if the corresponding condition was not met) and the resulting RowVersion. Note that Success property of the result only indicates successful completion as related to conditional Put operations and is always true for unconditional Puts. If the Put operation fails for any other reason, an exception will be thrown. Using fields of data type JSON allows more flexibility in the use of data as the data in JSON field does not have a predefined schema. To put value into a JSON field, supply a `MapValue` instance as its field value as part of the row value. You may also create its value from a JSON string via `FieldValue.FromJsonString`.

Download the full code ***AddData.cs*** from the examples here.

```
private static async Task insertData(NoSQLClient client, String acctdet){
   var putResult = await client.PutAsync(TableName,
                              FieldValue.FromJsonString(acctdet).AsMapValue);
   if (putResult.ConsumedCapacity != null)
   {
      Console.WriteLine(" Added a row to the stream_acct table");
   }
}
```

Inserting data into a JSON collection table: You add a row in the JSON collection table by putting its value from a JSON string created through `FromJsonString` method.

```
var tableName = "usersJSON";

private const string data= @"{
"id": 1,
"name": "John Doe",
"age": 25,
"college" : {"name" : "Presidency", "branch" : "Biotechnology"}
}

var result = await client.PutAsync(tableName,
FieldValue.FromJsonString(data).AsMapValue);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Using MultiWrite API to insert data

You can add more than a row of data in a single database operation using MultiWrite API.

- Java
- Python
- Go
- Node.js
- C#

Java

You can perform a sequence of `PutRequest` operations associated with a table that share the same shard key portion of their primary keys as a single atomic write operation using the `WriteMultipleRequest` class. You can also simultaneously add data to a parent and child table using the `WriteMultipleRequest` class. This is an efficient way to atomically modify multiple related rows. If the operation is successful, the `WriteMultipleResult.getSuccess()` method returns true.

See Oracle NoSQL Java SDK API Reference for more details on the various classes and methods.

Download the full code ***MultiWrite.java*** from the examples here.

```
private static void writeMul(NoSQLHandle handle,String parent_tblname,
String parent_data, String child_tblname, String child_data){
      WriteMultipleRequest umRequest = new WriteMultipleRequest();
      PutRequest putRequest =
         new
PutRequest().setValueFromJson(parent_data,null).setTableName(parent_tblname);
      umRequest.add(putRequest, false);
      putRequest =
         new
PutRequest().setValueFromJson(child_data,null).setTableName(child_tblname);
      umRequest.add(putRequest, false);
      WriteMultipleResult umResult = handle.writeMultiple(umRequest);
}
```

Python

You can perform a sequence of `PutRequest` operations associated with a table that share the same shard key portion of their primary keys as a single atomic write operation using the `borneo.WriteMultipleRequest` class. You can also simultaneously add data to a parent and child table using the `borneo.WriteMultipleRequest` class. This is an efficient way to atomically modify multiple related rows.

See Oracle NoSQL Python SDK API Reference for more details on the various classes and methods.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***MultiWrite.py*** from the examples here.

```
def mul_write(handle,parent_tblname,parent_data,
child_tblname, child_data):
    request = PutRequest()
    request.set_value_from_json(parent_data)
    request.set_table_name('ticket')
    wm_req.add(request, True)
    request1 = PutRequest()
    request1.set_table_name(child_tblname)
    request1.set_value_from_json(child_data)
    wm_req.add(request1, True)
    result = handle.write_multiple(wm_req)
```

Go

You can perform a sequence of `PutRequest` operations associated with a table that share the same shard key portion of their primary keys as a single atomic write operation using the `WriteMultipleRequest` class. You can also simultaneously add data to a parent and child table using the `WriteMultipleRequest` class. This is an efficient way to atomically modify multiple related rows.

See Oracle NoSQL Go SDK API Reference for more details on the various classes and methods.

Download the full code ***MultiWrite.go*** from the examples here.

```
//multiple write from the table
func mul_write(client *nosqldb.Client, err error, parenttbl_name string,
parent_data string, childtbl_name string, child_data string)(){
   value, err := types.NewMapValueFromJSON(parent_data)
    putReq := &nosqldb.PutRequest{
      TableName: parenttbl_name,
      Value:     value,
   }
    wmReq := &nosqldb.WriteMultipleRequest{
        TableName: "",
        Timeout:   10 * time.Second,
    }
   wmReq.AddPutRequest(putReq, true)

    value1, err := types.NewMapValueFromJSON(child_data)
    putReq1 := &nosqldb.PutRequest{
      TableName: childtbl_name,
      Value:     value1,
   }
   wmReq.AddPutRequest(putReq1, true)
    wmRes, err := client.WriteMultiple(wmReq)
    if err != nil {
        fmt.Printf("WriteMultiple() failed: %v\n", err)
        return
    }
    if wmRes.IsSuccess() {
        fmt.Printf("WriteMultiple() succeeded\n")
    } else {
        fmt.Printf("WriteMultiple() failed\n")
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}
```

Node.js

You can perform a sequence of put operations associated with a table that share the same shard key portion of their primary keys as a single atomic write operation using the `writeMany` method. You can also simultaneously add data to a parent and child table using the `writeMany` method. This is an efficient way to atomically modify multiple related rows.

For method details, see NoSQLClient class.

JavaScript: Download the full code ***MultiWrite.js*** from the examples here.

```
const ops = [
               {
                  tableName: 'ticket',
                  put: {
                     "ticketNo": "1762344493810",
                     "confNo" : "LE6J4Z"
                  },
                  abortOnFail: true
               },
               {
                  tableName: 'ticket.bagInfo',
                  put: {
                     "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"
                   },
                  abortOnFail: true
               }
            ];

const res = await handle.writeMany(ops, null);
```

TypeScript: Download the full code ***MultiWrite.ts*** from the examples here.

```
const ops = [
               {
                  tableName: 'ticket',
                  put: {
                     "ticketNo": "1762344493810",
                     "confNo" : "LE6J4Z"
                  },
                  abortOnFail: true
               },
               {
                  tableName: 'ticket.bagInfo',
                  put: {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                    "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"
                  },
                abortOnFail: true
            }
        ];

const res = await handle.writeMany(ops, null);
```

C#

You can perform a sequence of put operations associated with a table that share the same shard key portion of their primary keys as a single atomic write operation using the `PutManyAsync` method. You can also simultaneously add data to a parent and child table using the `PutManyAsync` method. This is an efficient way to atomically modify multiple related rows.

See Oracle NoSQL Dotnet SDK API Reference for more details of all classes and methods.

Download the full code ***MultiWrite.cs*** from the examples here.

```
private static async Task mul_write(NoSQLClient client,string parentbl_name,
string data1, string childtbl_name, string data2){
    var result = await client.WriteManyAsync(
        new WriteOperationCollection()
          .AddPut(parentbl_name, FieldValue.FromJsonString(data1).AsMapValue)
          .AddPut(childtbl_name, FieldValue.FromJsonString(data2).AsMapValue)
         );
    }
```

Upsert Data

The word `UPSERT` combines `UPDATE` and `INSERT`, describing the statement's function.

- Using SQL command to upsert data
- Using API to upsert data

Using SQL command to upsert data

Use an UPSERT statement to insert a row where it does not exist, or to update the row with new values when it does.

Example : Updating data in the `BaggageInfo` table using UPSERT command.

Copyright © 2022, 2025, Oracle and/or its affiliates.

The existing details for the customer with full name `Adam Phillips` is shown below.

```
SELECT * FROM BaggageInfo WHERE fullname="Adam Phillips"
```


```
{
  "ticketNo" : 1762344493810,
  "fullName" : "Adam Phillips",
  "gender" : "M",
  "contactPhone" : "893-324-1064",
  "confNo" : "LE6J4Z",
  "bagInfo" : [{
    "bagArrivalDate" : "2019-02-01T16:13:00Z",
    "flightLegs" : [{
      "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"
      }],
      "estimatedArrival" : "2019-02-01T11:00:00Z",
      "flightDate" : "2019-02-01T06:00:00Z",
      "flightNo" : "BM604",
      "fltRouteDest" : "LAX",
      "fltRouteSrc" : "MIA"
    }, {
      "actions" : [{
        "actionAt" : "MEL",
        "actionCode" : "Offload to Carousel at MEL",
        "actionTime" : "2019-02-01T16:15:00Z"
      }, {
        "actionAt" : "LAX",
        "actionCode" : "ONLOAD to MEL",
        "actionTime" : "2019-02-01T15:35:00Z"
      }, {
        "actionAt" : "LAX",
        "actionCode" : "OFFLOAD from LAX",
        "actionTime" : "2019-02-01T15:18:00Z"
      }],
      "estimatedArrival" : "2019-02-01T16:15:00Z",
      "flightDate" : "2019-02-01T06:13:00Z",
      "flightNo" : "BM667",
      "fltRouteDest" : "MEL",
      "fltRouteSrc" : "LAX"
    }],
    "id" : "79039899165297",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "MEL",
    "lastSeenTimeGmt" : "2019-02-01T16:13:00Z",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    "routing" : "MIA/LAX/MEL",
    "tagNum" : "17657806255240"
  }]
}
1 row returned
```

You modify the existing row using the `UPSERT` command. You can use an optional RETURNING clause to fetch the values after UPSERT is performed. The updated value for the customer with full name Adam Phillips is fetched as shown below.

```
UPSERT INTO BaggageInfo VALUES(
1762344493810,
"Adam Phillips",
"M",
"893-324-1864",
"LE6J4Y",
[ {
    "id" : "79039899165297",
    "tagNum" : "17657806255240",
    "routing" : "MIA/LAX/MEL",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "MEL",
    "flightLegs" : [ {
      "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"
      } ]
    }, {
      "flightNo" : "BM667",
      "flightDate" : "2019-02-01T06:13:00Z",
      "fltRouteSrc" : "LAX",
      "fltRouteDest" : "MEL",
      "estimatedArrival" : "2019-02-01T16:15:00Z",
      "actions" : [ {
        "actionAt" : "MEL",
        "actionCode" : "Offload to Carousel at MEL",
        "actionTime" : "2019-02-01T16:15:00Z"
      }, {
        "actionAt" : "LAX",
        "actionCode" : "ONLOAD to MEL",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            "actionTime" : "2019-02-01T15:35:00Z"
          }, {
            "actionAt" : "LAX",
            "actionCode" : "OFFLOAD from LAX",
            "actionTime" : "2019-02-01T15:18:00Z"
          } ]
        } ],
        "lastSeenTimeGmt" : "2019-02-01T16:18:00Z",
        "bagArrivalDate" : "2019-02-01T16:18:00Z"
      } ]
    ) RETURNING *


    {
      "ticketNo" : 1762344493810,
      "fullName" : "Adam Phillips",
      "gender" : "M",
      "contactPhone" : "893-324-1864",
      "confNo" : "LE6J4Y",
      "bagInfo" : [{
        "bagArrivalDate" : "2019-02-01T16:18:00Z",
        "flightLegs" : [{
          "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"
          }],
          "estimatedArrival" : "2019-02-01T11:00:00Z",
          "flightDate" : "2019-02-01T06:00:00Z",
          "flightNo" : "BM604",
          "fltRouteDest" : "LAX",
          "fltRouteSrc" : "MIA"
        }, {
          "actions" : [{
            "actionAt" : "MEL",
            "actionCode" : "Offload to Carousel at MEL",
            "actionTime" : "2019-02-01T16:15:00Z"
          }, {
            "actionAt" : "LAX",
            "actionCode" : "ONLOAD to MEL",
            "actionTime" : "2019-02-01T15:35:00Z"
          }, {
            "actionAt" : "LAX",
            "actionCode" : "OFFLOAD from LAX",
            "actionTime" : "2019-02-01T15:18:00Z"
          }],
          "estimatedArrival" : "2019-02-01T16:15:00Z",
          "flightDate" : "2019-02-01T06:13:00Z",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        "flightNo" : "BM667",
        "fltRouteDest" : "MEL",
        "fltRouteSrc" : "LAX"
      }],
      "id" : "79039899165297",
      "lastActionCode" : "OFFLOAD",
      "lastActionDesc" : "OFFLOAD",
      "lastSeenStation" : "MEL",
      "lastSeenTimeGmt" : "2019-02-01T16:18:00Z",
      "routing" : "MIA/LAX/MEL",
      "tagNum" : "17657806255240"
    }]
  }
```

> **Note**
>
> 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.

Example : Inserting data in the `BaggageInfo` table using UPSERT command.

A new entry value for a customer with full name `Birgit Naquin` is added using the `UPSERT` command.

```
SELECT * FROM BaggageInfo WHERE fullname="Birgit Naquin";

0 row returned


UPSERT INTO BaggageInfo VALUES(
1762392196147,
"Birgit Naquin",
"M",
"165-742-5715",
"QD1L0T",
[ {
    "id" : "7903989918469",
    "tagNum" : "17657806240229",
    "routing" : "JFK/MAD",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "MAD",
    "flightLegs" : [ {
      "flightNo" : "BM495",
      "flightDate" : "2019-03-07T07:00:00Z",
      "fltRouteSrc" : "JFK",
      "fltRouteDest" : "MAD",
      "estimatedArrival" : "2019-03-07T14:00:00Z",
      "actions" : [ {
        "actionAt" : "MAD",
        "actionCode" : "Offload to Carousel at MAD",
        "actionTime" : "2019-03-07T13:54:00Z"
      }, {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
          "actionAt" : "JFK",
          "actionCode" : "ONLOAD to MAD",
          "actionTime" : "2019-03-07T07:00:00Z"
        }, {
          "actionAt" : "JFK",
          "actionCode" : "BagTag Scan at JFK",
          "actionTime" : "2019-03-07T06:53:00Z"
        }, {
          "actionAt" : "JFK",
          "actionCode" : "Checkin at JFK",
          "actionTime" : "2019-03-07T05:03:00Z"
        } ]
      } ],
      "lastSeenTimeGmt" : "2019-03-07T13:51:00Z",
      "bagArrivalDate" : "2019-03-07T13:51:00Z"
    } ]
)

{"NumRowsInserted":1}

1 row returned
```

The result shows `{"NumRowsInserted":1}` which implies a new row has been inserted. The value inserted using the `UPSERT` command can be viewed as shown below:

```
SELECT * FROM BaggageInfo where fullname="Birgit Naquin"
{
  "ticketNo" : 1762392196147,
  "fullName" : "Birgit Naquin",
  "gender" : "M",
  "contactPhone" : "165-742-5715",
  "confNo" : "QD1L0T",
  "bagInfo" : [{
    "bagArrivalDate" : "2019-03-07T13:51:00Z",
    "flightLegs" : [{
      "actions" : [{
        "actionAt" : "MAD",
        "actionCode" : "Offload to Carousel at MAD",
        "actionTime" : "2019-03-07T13:54:00Z"
      }, {
        "actionAt" : "JFK",
        "actionCode" : "ONLOAD to MAD",
        "actionTime" : "2019-03-07T07:00:00Z"
      }, {
        "actionAt" : "JFK",
        "actionCode" : "BagTag Scan at JFK",
        "actionTime" : "2019-03-07T06:53:00Z"
      }, {
        "actionAt" : "JFK",
        "actionCode" : "Checkin at JFK",
        "actionTime" : "2019-03-07T05:03:00Z"
      }],
      "estimatedArrival" : "2019-03-07T14:00:00Z",
      "flightDate" : "2019-03-07T07:00:00Z",
      "flightNo" : "BM495",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        "fltRouteDest" : "MAD",
        "fltRouteSrc" : "JFK"
      }],
      "id" : "7903989918469",
      "lastActionCode" : "OFFLOAD",
      "lastActionDesc" : "OFFLOAD",
      "lastSeenStation" : "MAD",
      "lastSeenTimeGmt" : "2019-03-07T13:51:00Z",
      "routing" : "JFK/MAD",
      "tagNum" : "17657806240229"
    }]
}
1 row returned
```

> **Note**
>
> 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. You can also use an optional RETURNING clause as part of the UPSERT command.

Example : Use UPSERT statement to add/modify data in the `stream_acct` table.

```
UPSERT INTO stream_acct VALUES
(
   1,
   "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": 75,
               "minWatched": 75,
               "date": "2022-04-18"
            },
            {
               "episodeID": 30,
               "lengthMin": 60,
               "episodeName": "Season 1 episode 2",
               "minWatched": 40,
               "date": "2022 - 04 - 18 "
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                }]
            },
            {
                "seasonNum": 2,
                "numEpisodes": 2,
                "episodes": [{
                "episodeID": 40,
                    "episodeName": "Season 2 episode 1",
                    "lengthMin": 40,
                    "minWatched": 30,
                    "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": 20,
                    "minWatched": 20,
                    "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"
                },
                {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                "episodeID": 30,
                "episodeName": "Merci",
                "lengthMin": 42,
                "minWatched": 42,
                "date": "2022-03-08"
            }]
        }]
    }]
}
) RETURNING *
```

In the above example, a new row is inserted if the `stream_acct` table does not have a row corresponding to `acct_id =1`. Else the existing row with the value of `acct_id =1` is updated.

Example : 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 the JSON collection tables. Consider the JSON collection table created for a shopping application table.

```
UPSERT into storeAcct values ("1417114588", {"firstName" : "Dori",
"lastName" : "Martin", "email" : "dormartin@usmail.com", "address" :
{"Dropbox" : "Presidency College"}}) RETURNING *;
```

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"}
```

Using API to upsert data

You can use the UPSERT SQL command in the Query request to update or insert data.

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***ModifyData.java*** from the examples here.

```
/*Upsert data*/
private static void upsertRows(NoSQLHandle handle,String sqlstmt) throws
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)){
         for (MapValue res : results) {
            System.out.println("\t" + res);
         }
      }
}


String upsert_row = "UPSERT INTO stream_acct VALUES("+
         "1,"+
         "\"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\": 70,"+
                  "\"minWatched\": 70,"+
                  "\"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\": 40,"+
                  "\"minWatched\": 40,"+
                  "\"date\" : \"2022-04-25\""+
               "},"+
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                    "{"+
                        "\"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\""+
                    "}"+
                    "]"+
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            "}"+
            "]"+
         "}"+
         "]}) RETURNING *";
         System.out.println("Upsert data ");
         upsertRows(handle,upsert_row);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***ModifyData.py*** from the examples here.

```
#upsert data
def upsert_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   result = handle.query(request)
   print('Upsert data')
   for r in result.get_results():
      print('\t' + str(r))


upsert_row = '''
UPSERT INTO stream_acct VALUES
(
   1,
   "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": 75,
               "minWatched": 75,
               "date": "2022-04-18"
            },
            {
               "episodeID": 30,
               "lengthMin": 60,
               "episodeName": "Season 1 episode 2",
               "minWatched": 40,
               "date": "2022 - 04 - 18 "
            }]
         },
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        {
           "seasonNum": 2,
           "numEpisodes": 2,
           "episodes": [{
           "episodeID": 40,
              "episodeName": "Season 2 episode 1",
              "lengthMin": 40,
              "minWatched": 30,
              "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": 20,
              "minWatched": 20,
              "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",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                "lengthMin": 42,
                "minWatched": 42,
                "date": "2022-03-08"
            }]
        }]
    }]
  }
) RETURNING *
'''
upsert_data(handle,upsert_row)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***ModifyData.go*** from the examples here.

```
//upsert data in the table
func upsertRows(client *nosqldb.Client, err error,
                tableName string, querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
    }
    prepRes, err := client.Prepare(prepReq)
    if err != nil {
        fmt.Printf("Prepare failed: %v\n", err)
        return
    }
    queryReq := &nosqldb.QueryRequest{
        PreparedStatement: &prepRes.PreparedStatement,    }
    var results []*types.MapValue
    for {
        queryRes, err := client.Query(queryReq)
        if err != nil {
            fmt.Printf("Upsert failed: %v\n", err)
            return
        }
        res, err := queryRes.GetResults()
        if err != nil {
            fmt.Printf("GetResults() failed: %v\n", err)
            return
        }
        results = append(results, res...)
        if queryReq.IsDone() {
            break
        }
    }
    for i, r := range results {
        fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
    }
}


upsert_data := `UPSERT INTO stream_acct VALUES(
1,
"AP",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                "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": 75,
                                            "minWatched": 75,
                                            "date": "2022-04-18"
                                        },
                                        {
                                            "episodeID": 30,
                                            "lengthMin": 60,
                                            "episodeName": "Season 1 episode 2",
                                            "minWatched": 40,
                                            "date": "2022 - 04 - 18 "
                                        }
                                    ]
                                },
                                {
                                    "seasonNum": 2,
                                    "numEpisodes": 2,
                                    "episodes": [
                                        {
                                            "episodeID": 40,
                                            "episodeName": "Season 2 episode 1",
                                            "lengthMin": 40,
                                            "minWatched": 30,
                                            "date": "2022-04-25"
                                        },
                                        {
                                            "episodeID": 50,
                                            "episodeName": "Season 2 episode 2",
                                            "lengthMin": 45,
                                            "minWatched": 30,
                                            "date": "2022-04-27"
                                        }
                                    ]
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                },
                {
                    "seasonNum": 3,
                    "numEpisodes": 2,
                    "episodes": [
                        {
                            "episodeID": 60,
                            "episodeName": "Season 3 episode 1",
                            "lengthMin": 20,
                            "minWatched": 20,
                            "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"
                        }
                    ]
                }
            ]
        }
    ]
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
}) RETURNING *`

upsertRows(client, err,tableName,upsert_data)
```

Node.js

You can use the UPSERT SQL command in the Query request to update or insert data. To execute a query use `query` method.

JavaScript: Download the full code ***ModifyData.js*** from the examples here.

```
/*upserts data in the table*/
async function upsertData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

TypeScript: Download the full code ***ModifyData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


async function upsertData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


const upsert_row = `UPSERT INTO stream_acct VALUES
(
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
1,
"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": 75,
            "minWatched": 75,
            "date": "2022-04-18"
         },
         {
            "episodeID": 30,
            "lengthMin": 60,
            "episodeName": "Season 1 episode 2",
            "minWatched": 40,
            "date": "2022 - 04 - 18 "
         }]
      },
      {
         "seasonNum": 2,
         "numEpisodes": 2,
         "episodes": [{
            "episodeID": 40,
            "episodeName": "Season 2 episode 1",
            "lengthMin": 40,
            "minWatched": 30,
            "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": 20,
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                "minWatched": 20,
                "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"
             }]
          }]
       }]
    }) RETURNING *`

await upsertData(handle,upsert_row);
console.log("Upsert data into table");
```

C#

You can use the UPSERT SQL command in the Query request to update or insert data. To execute a query, you can use `QueryAsync` method or `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***ModifyData.cs*** from the examples here.

```
private static async Task upsertData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
  }
}


private const string upsert_row = @"UPSERT INTO stream_acct VALUES
(
   1,
   ""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"": 75,
               ""minWatched"": 75,
               ""date"": ""2022-04-18""
            },
            {
               ""episodeID"": 30,
               ""lengthMin"": 60,
               ""episodeName"": ""Season 1 episode 2"",
               ""minWatched"": 40,
               ""date"": ""2022 - 04 - 18""
            }]
         },
         {
            ""seasonNum"": 2,
            ""numEpisodes"": 2,
            ""episodes"": [{
               ""episodeID"": 40,
               ""episodeName"": ""Season 2 episode 1"",
               ""lengthMin"": 40,
               ""minWatched"": 30,
               ""date"": ""2022-04-25""
            },
            {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                ""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"": 20,
                ""minWatched"": 20,
                ""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""
             }]
          }]
       }]
    }
) RETURNING *";

await upsertData(client,upsert_row);
Console.WriteLine("Upsert data in table");
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Update Data

- Using SQL command to update data
- Using API to update data

Using SQL command to update data

An update statement can be used to update a row in a table. It now 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.

- The update statement uses update clauses to modify the values of one or more fields in a table. Oracle NoSQL Database supports the following update clauses:
  - SET clause updates the value of one or more existing fields.
  - ADD clause adds new elements in one or more arrays.
  - PUT clause adds new fields in one or more maps. It can also update the values of existing map fields.
  - REMOVE clause removes elements/fields from one or more arrays/maps.
  - JSON MERGE clause allows you to make a set of changes to a JSON document.
  - SET TTL clause updates the expiration time of the row.
- The `WHERE` clause specifies what row to update. The current implementation supports both single-row and multiple-row updates, so the `WHERE` clause must specify a complete primary key for single-row updates and for multiple-row updates you must specify a shard key, ensuring that all rows reside on the same shard. This enables Oracle NoSQL Database to perform this update in an ACID transaction.
- There is an optional `RETURNING` clause which 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. The `RETURNING` clause works only when the `WHERE` clause includes the complete primary key.

> **Note**
>
> Currently, it does not support multiple-row updates.

- Furthermore, if no row satisfies the `WHERE` conditions, the update statement returns an empty result.
- There is also a limit on the number of records a single update query can update. By default, an update query can modify up to 1,000 records. To override it, use `setLimit(int limit)` method, but avoid setting it too high, as it may cause timeouts or increased latency. See `QueryRequest.setLimit(int limit)`, for more details.

Example 1: Simple example to change the column values.

```
UPDATE BaggageInfo
SET contactPhone = "823-384-1964",
confNo = "LE6J4Y"
WHERE ticketNo = 1762344493810
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Explanation: In the query above, you use the SET clause to update a few column values for a given ticket number.

Example 2: Update row data and fetch the values with a RETURNING clause.

```
UPDATE BaggageInfo
SET contactPhone = "823-384-1964",
confNo = "LE6J4Y"
WHERE ticketNo = 1762344493810 RETURNING *
```

Explanation: In the query above, you use the RETURNING clause to fetch back the data after the UPDATE transaction is completed.

Output:


```
{"ticketNo":1762344493810,"fullName":"Adam
Phillips","gender":"M","contactPhone":"823-384-1964",
"confNo":"LE6J4Y",
"bagInfo":{"bagInfo":[{"bagArrivalDate":"2019.02.02 at 03:13:00
AEDT","flightLegs":
[{"actions":[{"actionAt":"MIA","actionCode":"ONLOAD to
LAX","actionTime":"2019.02.01 at 01:13:00 EST"},
{"actionAt":"MIA","actionCode":"BagTag Scan at MIA","actionTime":"2019.02.01
at 00:47:00 EST"},
{"actionAt":"MIA","actionCode":"Checkin at MIA","actionTime":"2019.01.31 at
23:38:00 EST"}],
"estimatedArrival":"2019.02.01 at 03:00:00 PST","flightDate":"2019.02.01 at
01:00:00 EST",
"flightNo":"BM604","fltRouteDest":"LAX","fltRouteSrc":"MIA"},{"actions":
[{"actionAt":"MEL","actionCode":"Offload to Carousel at
MEL","actionTime":"2019.02.02 at 03:15:00 AEDT"},
{"actionAt":"LAX","actionCode":"ONLOAD to MEL","actionTime":"2019.02.01 at
07:35:00 PST"},
{"actionAt":"LAX","actionCode":"OFFLOAD from LAX","actionTime":"2019.02.01 at
07:18:00 PST"}],
"estimatedArrival":"2019.02.02 at 03:15:00 AEDT","flightDate":"2019.01.31 at
22:13:00 PST",
"flightNo":"BM667","fltRouteDest":"MEL","fltRouteSrc":"LAX"}],"id":"7903989916
5297",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MEL",
"lastSeenTimeGmt":"2019.02.02 at 03:13:00 AEDT","routing":"MIA/LAX/
MEL","tagNum":"17657806255240"}]}}
```


Example 3: Update the account expiry date for a customer in the `stream_acct` table.

```
UPDATE stream_acct SET account_expiry="2023-12-28T00:00:00.0Z" WHERE acct_Id=3
```

Explanation: In the query above, you use the SET clause to update the `account_expiry` field for a given account in the streaming media application.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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"}})
```

Example 1: Use `UPDATE` statement to update multiple rows

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 `info.address.city` field to 'Oslo', the `info.address.country` field to 'Norway', and increases the `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:

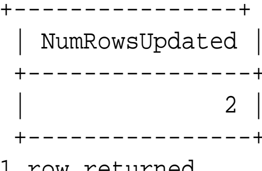

```
+-----------------+
 | NumRowsUpdated |
 +-----------------+
 |              2 |
 +-----------------+
1 row returned
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Now, run the `SELECT` query to verify the updated rows. Check that you have updated the `info.address.city` and `info.address.country` fields to 'Oslo' and 'Norway', respectively, and confirm that you have increased the `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      |
 +-------+------------+-----------------+---------------------------+
 |   101 | HR         | Liam Phillips   | address                   |
 |       |            |                 |     city    | Oslo        |
 |       |            |                 |     country | Norway      |
 |       |            |                 | salary      | 105000      |
 +-------+------------+-----------------+---------------------------+
 |   102 | HR         | Emma Johnson    | address                   |
 |       |            |                 |     city    | Oslo        |
 |       |            |                 |     country | Norway      |
 |       |            |                 | salary      | 110000      |
 +-------+------------+-----------------+---------------------------+
 |   104 | Finance    | Sophia Becker   | address                   |
 |       |            |                 |     city    | Munich      |
 |       |            |                 |     country | Germany     |
 |       |            |                 | salary      | 130000      |
 +-------+------------+-----------------+---------------------------+

4 rows returned
```

Example 2: Use `UPDATE` statement with a `REURNING` clause

Example 2a: To update a single row with a `RETURNING` clause

```
UPDATE EmployeeInfo emp SET emp.info.salary = emp.info.salary + 1000 WHERE
empID=101 and department="HR" RETURNING *
```

Explanation: In the above query, the `SET` clause increments the salary by 1000 for an employee working in HR department with an employee ID of 101, followed by the `RETURNING` clause. Since the `WHERE` clause specifies the complete primary key, the update affects a single row, and the output is returned directly due to the presence of the `RETURNING` clause.

Output:

```
+--------+------------+--------------+---------------------+
 | empID | department |   fullName    |         info        |
 +-------+------------+---------------+---------------------+
 |   101 | HR         | Liam Phillips | address             |
 |       |            |               |     city    | Oslo  |
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
 |       |           |              |     country | Norway |
 |       |           |              | salary      | 106000 |
 +-------+-----------+--------------+---------------------+

1 row returned
```

Example 2b: To update multiple rows with a `RETURNING` clause

```
UPDATE EmployeeInfo emp SET emp.info.salary = emp.info.salary + 1000 WHERE
department="HR" RETURNING *
```

Explanation: In the above query, the `SET` clause guides to increment the salary by 1000 for employees in HR department, followed by the `RETURNING` clause. As there are multiple employees in the HR department, the query throws an error because the `RETURNING` clause is currently not supported for multiple-row updates, and the `WHERE` clause must specify the complete primary key.

Output:

```
Error handling command UPDATE EmployeeInfo emp SET emp.info.salary =
emp.info.salary + 1000 WHERE department="HR" RETURNING *: Error: RETURNING
clause is not supported unless the complete primary key is specified in the
WHERE clause.
```

Updating JSON data

Example 1: Update a few fields, insert a nested JSON document, and remove an existing document from the `bagInfo` JSON array.

```
UPDATE BaggageInfo b
JSON MERGE b.bagInfo[0].flightLegs[size(b.bagInfo[0].flightLegs)-1] WITH
PATCH {"flightNo" : "BM107", "actions" : NULL, "tempActions" : {
        "actionAt" : "LAX",
        "actionStatus" : "in transit"
      }},
WHERE ticketNo = 1762344493810 RETURNING *
```

Explanation: The query above updates a passenger's baggage tracking information in an airline application. The `bagInfo` field is a JSON array that stores the passenger's baggage tracking information. You use the JSON MERGE patch to update the values in the last leg of the passenger's travel itinerary. To calculate the last element of the `flightLegs` JSON array, you use the built-in size function, which returns the number of travel segments and subtract it by one. In the patch expression, you supply the fields to be updated. Here, you modify the `flightNo` value, remove the actions field by supplying a NULL value, and insert a new JSON document `tempActions` in the `flightLegs` JSON array.

Output:

```
{
  "ticketNo" : 1762344493810,
  "fullName" : "Adam Phillips",
  "gender" : "M",
  "contactPhone" : "893-324-1064",
  "confNo" : "LE6J4Z",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
  "bagInfo" : [{
    "bagArrivalDate" : "2019-02-01T16:13:00Z",
    "flightLegs" : [{
      "actions" : [{ ... }],
      "estimatedArrival" : "2019-02-01T11:00:00Z",
      "flightDate" : "2019-02-01T06:00:00Z",
      "flightNo" : "BM604",
      "fltRouteDest" : "LAX",
      "fltRouteSrc" : "MIA"
    }, {
      "estimatedArrival" : "2019-02-01T16:15:00Z",
      "flightDate" : "2019-02-01T06:13:00Z",
      "flightNo" : "BM107",
      "fltRouteDest" : "MEL",
      "fltRouteSrc" : "LAX",
      "tempActions" : {
        "actionAt" : "LAX",
        "actionStatus" : "in transit"
      }
    }],
    "id" : "79039899165297",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "MEL",
    "lastSeenTimeGmt" : "2019-02-01T16:13:00Z",
    "routing" : "MIA/LAX/MEL",
    "tagNum" : "17657806255240"
  }]
}
```

Note

Here, the `bagInfo` field 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` JSON 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.

Example 2: Use JSON MERGE patch and PUT clause in an update statement.

You can use the JSON MERGE patch with all other update clauses (SET, ADD, PUT, REMOVE) in a single UPDATE statement as follows:

```
UPDATE BaggageInfo b
JSON MERGE b.bagInfo[0].flightLegs WITH PATCH {"flightNo" : "BM107",
"actions" : NULL, "tempActions" : {
        "actionAt" : "LAX",
        "actionStatus" : "in transit"
      }},
JSON MERGE b.bagInfo[1].flightLegs WITH PATCH {"flightNo" : "BM107",
"actions" : NULL, "tempActions" : {
        "actionAt" : "LAX",
        "actionStatus" : "in transit"
      }},
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
PUT b.bagInfo[0]{"bagArrivalDate" : "2019-03-13T00:00:00Z",
"lastSeenStation" : "SFO", "routing" : "SFO/LAX",},
PUT b.bagInfo[1]{"bagArrivalDate" : "2019-03-13T00:00:00Z",
"lastSeenStation" : "SFO", "routing" : "SFO/LAX",}
WHERE ticketNo = 1762320369957 RETURNING *
```

Explanation: In the query above, the passenger's record includes two checked bags. You use the JSON MERGE clause iteratively to update the baggage tracking information in both the elements of the `bagInfo` JSON array. Notice that the path expression in the target includes `flightLegs` object instead of an individual element of the `flightLegs` array. Therefore, the patch replaces the entire `flightLegs` JSON array with the `flightLegs` JSON object.

The PUT clause updates the `bagArrivalDate`, `lastSeenStation`, and `routing` field values in the specified elements of the `bagInfo` JSON array.

Output:

```
{
  "ticketNo" : 1762320369957,
  "fullName" : "Lorenzo Phil",
  "gender" : "M",
  "contactPhone" : "364-610-4444",
  "confNo" : "QI3V6Z",
  "bagInfo" : [{
    "bagArrivalDate" : "2019-03-13T00:00:00Z",
    "flightLegs" : {
      "flightNo" : "BM107",
      "tempActions" : {
        "actionAt" : "LAX",
        "actionStatus" : "in transit"
      }
    },
    "id" : "79039899187755",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "SFO",
    "lastSeenTimeGmt" : "2019-03-12T15:05:00Z",
    "routing" : "SFO/LAX",
    "tagNum" : "17657806240001"
  }, {
    "bagArrivalDate" : "2019-03-13T00:00:00Z",
    "flightLegs" : {
      "flightNo" : "BM107",
      "tempActions" : {
        "actionAt" : "LAX",
        "actionStatus" : "in transit"
      }
    },
    "id" : "79039899197755",
    "lastActionCode" : "OFFLOAD",
    "lastActionDesc" : "OFFLOAD",
    "lastSeenStation" : "SFO",
    "lastSeenTimeGmt" : "2019-03-12T15:05:00Z",
    "routing" : "SFO/LAX",
    "tagNum" : "17657806340001"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
  }]
}
```

Updating rows in a JSON collection table

Consider a row from the JSON collection table created for a shopping application.

Example: Modify the erroneous shopper data record in the `storeAcct` table.

You can use the UPDATE statement to update fields in an existing document in the JSON collection tables. The UPDATE operation works in the same way as fixed schema tables.

```
{"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"}]}
```

Use the 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:

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","orderID":"101200,"priceperunit":950,"status":"Preparing to dispatch"},
{"EstDelivery":"2023-11-01","item":"Wall
paper","orderID":"101200,"priceperunit":950,"status":"Transit"}]}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Using API to update data

You can use the UPDATE SQL command in the Query request to update data.

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***ModifyData.java*** from the examples here.

```
//Update data
private static void updateRows(NoSQLHandle handle,String sqlstmt) throws
Exception {
   QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
   handle.query(queryRequest);
   System.out.println("Updated table " + tableName);
}


/* update non-JSON data*/
String upd_stmt ="UPDATE stream_acct SET
account_expiry=\"2023-12-28T00:00:00.0Z\" WHERE acct_Id=3";
updateRows(handle,upd_stmt);
```

> **Note**
>
> Oracle NoSQL Database also supports user-defined row metadata. You can annotate actual row data with additional information in a JSON string using the `setRowMetadata()` method. To learn how it works, see Using row metadata in Write Operations.

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***ModifyData.py*** from the examples here.

```
#update data
def update_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   result = handle.query(request)
   print('Data Updated in table: stream_acct')


# update non-JSON data
upd_stmt ='''UPDATE stream_acct SET account_expiry="2023-12-28T00:00:00.0Z"
WHERE acct_Id=3'''
update_data(handle,upd_stmt)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***ModifyData.go*** from the examples here.

```
//update data in the table
func updateRows(client *nosqldb.Client, err error, tableName string,
querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
   queryReq := &nosqldb.QueryRequest{
      PreparedStatement: &prepRes.PreparedStatement,   }
      var results []*types.MapValue
      for {
         queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Upsert failed: %v\n", err)
         return
      }
      res, err := queryRes.GetResults()
      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
         return
      }
      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
     }
     for i, r := range results {
        fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
     }
     fmt.Printf("Updated data in the table: \n")
}


updt_stmt := "UPDATE stream_acct SET account_expiry='2023-12-28T00:00:00.0Z'
WHERE acct_Id=3"
updateRows(client, err,tableName,updt_stmt)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***ModifyData.js*** from the examples here.

```
/*updates data in the table*/
async function updateData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

TypeScript: Download the full code ***ModifyData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


async function updateData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


const updt_stmt = 'UPDATE stream_acct SET
account_expiry="2023-12-28T00:00:00.0Z" WHERE acct_Id=3'

await updateData(handle,updt_stmt);
console.log("Data updated in the table");
```

C#

You can use the UPDATE SQL command in the Query request to update data. To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***ModifyData.cs*** from the examples here.

```
private static async Task updateData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
}


private const string updt_stmt =
@"UPDATE stream_acct SET account_expiry =""2023-12-28T00:00:00.0Z"" WHERE
acct_Id=3";

await updateData(client,updt_stmt);
Console.WriteLine("Data updated in the table");
```

Modify JSON data

- Using SQL command
- Using API

Using SQL command

While updating JSON data, in addition to `WHERE`, `SET` and `RETURNING` clause, the following clauses can be used..

- 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 `PUT` clause is used primarily to add new fields to a JSON document. It consists of a target expression, which should normally return one or more fields to be inserted into the target JSON document.
- The `REMOVE` clause consists of a single target expression, which computes the items to be removed.

Example 1: Update table and add data in a JSON object

Add elements to the action array (at a given array element) for a particular flight Leg of a passenger. By default, the element is added at the end. If a number is specified, it is inserted in that position. In the example below, you want the new element to be added in the 2nd position.

```
UPDATE BaggageInfo bag
ADD bag.bagInfo[0].flightLegs[0].actions 2  {"actionAt" : "LAX",
         "actionCode" : "WAITING at LAX",
         "actionTime" : "2019-02-01T06:13:00Z"}
WHERE ticketNo=1762344493810
RETURNING *
```

Example 2: Update table and update data from a JSON object.

Copyright © 2022, 2025, Oracle and/or its affiliates.

You could update the data from a JSON object using the SET clause. Here the second element of the actions array is updated with new values for a given ticket number.

```
UPDATE BaggageInfo bag
SET bag.bagInfo[0].flightLegs[0].actions[2]=
{"actionAt" : "LAX",
"actionCode" : "STILL WAITING at LAX",
"actionTime" : "2019-02-01T06:15:00Z"}
WHERE ticketNo=1762344493810 RETURNING *
```

Example 3: Update table and remove data in a JSON object.

You can use the `REMOVE` clause to remove a given element from an array. You need to specify which element of the array needs to be removed using the index of the array.

```
UPDATE BaggageInfo bag
REMOVE bag.bagInfo[0].flightLegs[0].actions[1]
WHERE ticketNo=1762344493810
RETURNING *
```

Example 4: Update stream_acct table and add and remove data in a JSON object.

In the `stream_acct` table, for a customer you can add the details of a particular series episode of a show using the ADD clause in the UPDATE statement.

```
UPDATE stream_acct acct1 ADD
acct1.acct_data.contentStreamed.seriesInfo[1].episodes {
     "date" : "2022-04-26",
     "episodeID" : 43,
     "episodeName" : "Season 2 episode 2",
     "lengthMin" : 45,
     "minWatched" : 45} WHERE acct_Id=2 RETURNING *
```

Similarly , you can remove the details of a particular series episode of a show using the REMOVE clause in the UPDATE statement.

```
UPDATE stream_acct acct1
REMOVE acct1.acct_data.contentStreamed.seriesInfo[1].episodes[1]
WHERE acct_Id=2 RETURNING *
```

Using API

You can use the UPDATE SQL command to add and remove data in a JSON object in your table.

- Java
- Python
- Go
- Node.js

Copyright © 2022, 2025, Oracle and/or its affiliates.

- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***ModifyData.java*** from the examples here.

```
//Update data
private static void updateRows(NoSQLHandle handle,String sqlstmt) throws
Exception {
   QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
   handle.query(queryRequest);
   System.out.println("Updated table " + tableName);
}


/* update JSON data and add a node*/
String upd_json_addnode="UPDATE stream_acct acct1 ADD
acct1.acct_data.contentStreamed.seriesInfo[1].episodes
   {\"date\" : \"2022-04-26\","+
   "\"episodeID\" : 43,"+
   "\"episodeName\" : \"Season 2 episode 2\","+
   "\"lengthMin\" : 45,"+
   "\"minWatched\" : 45} WHERE acct_Id=2 RETURNING *";
updateRows(handle,upd_json_addnode);

/* update JSON data and remove a node*/
String upd_json_delnode="UPDATE stream_acct acct1 REMOVE
acct1.acct_data.contentStreamed.seriesInfo[1].episodes[1] WHERE acct_Id=2
RETURNING *";
updateRows(handle,upd_json_delnode);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***ModifyData.py*** from the examples here.

```
#update data
def update_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   result = handle.query(request)
   print('Data Updated in table: stream_acct')


# update JSON data and add a node
upd_json_addnode = '''UPDATE stream_acct acct1 ADD
acct1.acct_data.contentStreamed.seriesInfo[1].episodes {
     "date" : "2022-04-26",
     "episodeID" : 43,
     "episodeName" : "Season 2 episode 2",
     "lengthMin" : 45,
     "minWatched" : 45} WHERE acct_Id=2 RETURNING *'''
update_data(handle,upd_json_addnode)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
# update JSON data and delete a node
upd_json_delnode = '''UPDATE stream_acct acct1 REMOVE
acct1.acct_data.contentStreamed.seriesInfo[1].episodes[1] WHERE acct_Id=2
RETURNING *'''
update_data(handle,upd_json_delnode)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***ModifyData.go*** from the examples here.

```
//update data in the table
func updateRows(client *nosqldb.Client, err error, tableName string,
querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
   queryReq := &nosqldb.QueryRequest{
      PreparedStatement: &prepRes.PreparedStatement,    }
      var results []*types.MapValue
      for {
         queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Upsert failed: %v\n", err)
         return
      }
      res, err := queryRes.GetResults()
      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
         return
      }
      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
     }
     for i, r := range results {
        fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
     }
     fmt.Printf("Updated data in the table: \n")
}


upd_json_addnode := `UPDATE stream_acct acct1 ADD
acct1.acct_data.contentStreamed.seriesInfo[1].episodes {
       "date" : "2022-04-26",
       "episodeID" : 43,
       "episodeName" : "Season 2 episode 2",
       "lengthMin" : 45,
       "minWatched" : 45} WHERE acct_Id=2 RETURNING *`
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
updateRows(client, err,tableName,upd_json_addnode)

upd_json_delnode := `UPDATE stream_acct acct1 REMOVE
acct1.acct_data.contentStreamed.seriesInfo[1].episodes[1]
WHERE acct_Id=2 RETURNING *`
updateRows(client, err,tableName,upd_json_delnode)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***ModifyData.js*** from the examples here.

```
/*updates data in the table*/
async function updateData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

TypeScript: Download the full code ***ModifyData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


async function updateData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


const upd_json_addnode =
`UPDATE stream_acct acct1 ADD
acct1.acct_data.contentStreamed.seriesInfo[1].episodes {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   "date" : "2022-04-26",
   "episodeID" : 43,
   "episodeName" : "Season 2 episode 2",
   "lengthMin" : 45,
   "minWatched" : 45} WHERE acct_Id=2 RETURNING *`
await updateData(handle,upd_json_addnode);
console.log("New data node added in the table");

const upd_json_delnode =
'UPDATE stream_acct acct1 REMOVE
acct1.acct_data.contentStreamed.seriesInfo[1].episodes[1]
WHERE acct_Id=2 RETURNING *'
await updateData(handle,upd_json_delnode);
console.log("New Data node removed from the table");
```

C#

You can use the UPDATE SQL command to add and remove data in a JSON object in your table. To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***ModifyData.cs*** from the examples here.

```
private static async Task updateData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
}


private const string upd_json_addnode =
@"UPDATE stream_acct acct1 ADD
acct1.acct_data.contentStreamed.seriesInfo[1].episodes {
        ""date"" : ""2022-04-26"",
        ""episodeID"" : 43,
        ""episodeName"" : ""Season 2 episode 2"",
        ""lengthMin"" : 45,
        ""minWatched"" : 45} WHERE acct_Id=2 RETURNING *";

await updateData(client,upd_json_addnode);
Console.WriteLine("New data node added in the table");

private const string upd_json_delnode =
"UPDATE stream_acct acct1 REMOVE
acct1.acct_data.contentStreamed.seriesInfo[1].episodes[1]
WHERE acct_Id=2 RETURNING *";
await updateData(client,upd_json_delnode);
Console.WriteLine("New Data node removed from the table");
```

Delete Data

- Using SQL command to delete data
- Using API to delete a single row
- Using API to delete multiple rows

Copyright © 2022, 2025, Oracle and/or its affiliates.

- Using Query API to delete data

Using SQL command to delete data

The DELETE statement is used to remove 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.

Example 1: Delete data from a table with a simple WHERE clause.

You delete the data corresponding to a user with a given fullname.

```
DELETE FROM BaggageInfo
WHERE fullName = "Bonnie Williams"
```

Example 2: Delete data from a table with a RETURNING clause.

The RETURNING clause fetches the details of the row to be deleted. In the example below, you are fetching the full name and conf number corresponding to a ticket number which will be deleted.

```
DELETE FROM BaggageInfo
WHERE ticketNo = 1762392196147
RETURNING fullName,confNo
```

Output:

```
{"fullName":"Birgit Naquin","confNo":"QD1L0T"}
```

> **Note**
>
> 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 coordinated using a two-phase commit and lock mechanism. In such cases, it is recommended that the application re-run the DELETE statement.

Example 3: Delete data from `stream_acct` table based on the last name.

```
DELETE FROM stream_acct acct1
WHERE acct1.acct_data.firstName="Adelaide"
 AND acct1.acct_data.lastName="Willard"
```

Using API to delete a single row

You can use the `DeleteRequest` API and delete a single row using a primary key.

Copyright © 2022, 2025, Oracle and/or its affiliates.

The `DeleteRequest` API can be used to perform unconditional and conditional deletes.

- Delete any existing row. This is the default.
- Succeed only if the row exists and its version matches a specific version. Use `setMatchVersion` for this case.

- Java
- Python
- Go
- Node.js
- C#

Java

Download the full code ***ModifyData.java*** from the examples here.

```
//delete row based on primary KEY
private static void delRow(NoSQLHandle handle, MapValue m1) throws Exception {
   DeleteRequest delRequest = new
DeleteRequest().setKey(m1).setTableName(tableName);
   DeleteResult del = handle.delete(delRequest);
   if (del.getSuccess()) {
     System.out.println("Delete succeed");
   }
   else {
     System.out.println("Delete failed");
   }
}


/*delete a single row*/
MapValue m1= new MapValue();
m1.put("acct_Id",1);
delRow(handle,m1);
```

> **Note**
>
> Oracle NoSQL Database supports user-defined row metadata. You can annotate actual row data with additional information in a JSON string using the `setRowMetadata()` method. To understand more about the row metadata when you delete a row, see Using row metadata in Write Operations.

Python

Single rows are deleted using `borneo.DeleteRequest` using a primary key value.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***ModifyData.py*** from the examples here.

```
#del row with a primary KEY
def del_row(handle,table_name):
   request = DeleteRequest().set_key({'acct_Id':
1}).set_table_name(table_name)
   result = handle.delete(request)
   print('Deleted data from table: stream_acct')


# delete row based on primary key
del_row(handle,'stream_acct')
```

Go

The `DeleteRequest` is used to delete a row from a table. The row is identified using a primary key specified in `DeleteRequest.Key`.

Download the full code ***ModifyData.go*** from the examples here.

```
//delete with primary key
func delRow(client *nosqldb.Client, err error, tableName string)(){
   key := &types.MapValue{}
   key.Put("acct_Id",1)
   delReq := &nosqldb.DeleteRequest{
       TableName: tableName,
       Key:       key,
   }
   delRes, err := client.Delete(delReq)
   if err != nil {
    fmt.Printf("failed to delete a row: %v", err)
    return
   }
   if delRes.Success {
    fmt.Println("Delete succeeded")
   }
}


delRow(client, err,tableName)
```

Node.js

Use the `delete` method to delete a row from a table. For method details, see NoSQLClient class.

You must pass the table name and primary key of the row. In addition, you can make the delete operation conditional by specifying a `RowVersion` of the row that was previously returned by get or put methods.

JavaScript: Download the full code ***ModifyData.js*** from the examples here.

```
/*delete row based on primary key*/
async function delRow(handle) {
   try {
      /* Unconditional delete, should succeed.*/
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
      var result = await handle.delete(TABLE_NAME, { acct_Id: 1 });
      /* Expected output: delete succeeded*/
      console.log('delete ' + result.success ? 'succeeded' : 'failed');
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


await delRow(handle);
console.log("Row deleted based on primary key");
```

TypeScript: Download the full code ***ModifyData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/*delete row based on primary key*/
async function delRow(handle: NoSQLClient) {
   try {
      /* Unconditional delete, should succeed.*/
      var result = await handle.delete<StreamInt>(TABLE_NAME, { acct_Id: 1 });
      /* Expected output: delete succeeded*/
      console.log('delete ' + result.success ? 'succeeded' : 'failed');
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


await delRow(handle);
console.log("Row deleted based on primary key");
```

C#

To delete a row, use `DeleteAsync` method. Pass to it the table name and primary key of the row to delete. This method takes the primary key as MapValue. The field names should be the same as the table primary key column names.

`DeleteAsync` and `DeleteIfVersionAsync` methods return `Task<DeleteResult<RecordValue>>`. DeleteResult instance contains success status of the Delete operation. Delete operation may fail if the row with given primary key does not exist or this is a conditional Delete and provided row version did not match the existing row version.

Download the full code ***ModifyData.cs*** from the examples here.

```
private static async Task delRow(NoSQLClient client){
   var primaryKey = new MapValue
   {
      ["acct_Id"] = 1
   };
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   // Unconditional delete, should succeed.
   var deleteResult = await client.DeleteAsync(TableName, primaryKey);
   // Expected output: Delete succeeded.
   Console.WriteLine("Delete {0}.",deleteResult.Success ? "succeeded" :
"failed");
}


await delRow(client);
Console.WriteLine("Row deleted based on primary key");
```

Using API to delete multiple rows

You can use the `MultiDeleteRequest` API and delete more than one row from a NoSQL table.

You can use `MultiDeleteRequest` to delete multiple rows from a table in an atomic operation. The key used may be partial but must contain all of the fields that are in the shard key. A range may be specified to delete a range of keys. As this operation can exceed the maximum amount of data that can be modified in a single operation, a continuation key can be used to continue the operation.

- Java
- Python
- Go
- Node.js
- C#

Java

If a table's primary key is <YYYYMM, timestamp> and the its shard key is the YYYYMM, then all records that hit in the same month would be in same shard. It is possible to delete a range of timestamp values for a specific month using `MultiDeleteRequest` class.

See Oracle NoSQL Java SDK API Reference for more details on the various classes and methods.

Download the full code ***MultiDataOps.java*** from the examples here.

```
//Delete multiple rows from the table
private static void delMulRows(NoSQLHandle handle,int pinval) throws
Exception {
   MapValue key = new MapValue().put("pin", 1234567);
   MultiDeleteRequest multiDelRequest = new MultiDeleteRequest()
          .setKey(key)
          .setTableName(tableName);

   MultiDeleteResult mRes = handle.multiDelete(multiDelRequest);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   System.out.println("MultiDelete result = " + mRes);
}


/*delete multiple rows using shard key*/
delMulRows(handle,1234567);
```

> **Note**
>
> Oracle NoSQL Database supports user-defined row metadata. You can annotate actual row data with additional information in a JSON string using the `setRowMetadata()` method. To understand more about the row metadata when you delete multiple rows, see Using row metadata in Write Operations.

Python

You can use `borneo.MultiDeleteRequest` class to perform multiple deletes in a single atomic operation.

See Oracle NoSQL Python SDK API Reference for more details on the various classes and methods.

Download the full code ***MultiDataOps.py*** from the examples here.

```
#delete multiple rows
def multirow_delete(handle,table_name,pinval):
   request = MultiDeleteRequest().set_table_name(table_name).set_key({'pin':
pinval})
   result = handle.multi_delete(request)
)


/*delete multiple rows using shard key*/
multirow_delete(handle,'examplesAddress',1234567)
```

Go

You can use `MultiDelete` method to delete multiple rows from a table in a single atomic operation.

See Oracle NoSQL Go SDK API Reference for more details on the various classes and methods.

Download the full code ***MultiDataOps.go*** from the examples here.

```
//delete multiple rows
func delMulRows(client *nosqldb.Client, err error, tableName string,pinval
int)(){
   shardKey := &types.MapValue{}
   shardKey.Put("pin", pinval)
   multiDelReq := &nosqldb.MultiDeleteRequest{
        TableName: tableName,
        Key:       shardKey,
   }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   multiDelRes, err := client.MultiDelete(multiDelReq)
   if err != nil {
      fmt.Printf("failed to delete multiple rows: %v", err)
      return
   }
   fmt.Printf("MultiDelete result=%v\n", multiDelRes)
}


/*delete multiple rows using shard key*/
delMulRows(client, err,tableName,1234567)
```

Node.js

You can delete multiple rows having the same shard key in a single atomic operation using the `deleteRange` method.

JavaScript: Download the full code ***MultiDataOps.js*** from the examples here.

```
//deletes multiple rows
async function mulRowDel(handle,pinval){
   try {
      /* Unconditional delete, should succeed.*/
      var result = await handle.deleteRange(TABLE_NAME, { pin: pinval });
      /* Expected output: delete succeeded*/
      console.log('delete ' + result.success ? 'succeeded' : 'failed');
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


/*delete multiple rows using shard key*/
await mulRowDel(handle,1234567);
```

TypeScript: Download the full code ***MultiDataOps.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


//deletes multiple rows
async function mulRowDel(handle: NoSQLClient,pinVal: Integer){
   try {
      /* Unconditional delete, should succeed.*/
      var result = await handle.deleteRange<StreamInt>(TABLE_NAME, { pin: 
pinval });
      /* Expected output: delete succeeded*/
      console.log('delete ' + result.success ? 'succeeded' : 'failed');
   } catch(error) {
      console.error('  Error: ' + error.message);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}


/*delete multiple rows using shard key*/
await mulRowDel(handle,1234567);
```

C#

You can delete multiple rows having the same shard key in a single atomic operation using `DeleteRangeAsync` method.

Download the full code ***MultiDataOps.cs*** from the examples here.

```
//delete multiple rows
private static async Task mulDelRows(NoSQLClient client,int pinval){
   var parKey = new MapValue {["pin"] = pinval};
   var options = new DeleteRangeOptions();
   do
   {
      var result = await client.DeleteRangeAsync(TableName,parKey,options);
      Console.WriteLine($"Deleted {result.DeletedCount} row(s)");
      options.ContinuationKey = result.ContinuationKey;
   } while(options.ContinuationKey != null);
}


/*delete multiple rows using shard key*/
await mulDelRows(client,1234567);
```

Using Query API to delete data

You can use the `QueryRequest` API and delete one or more rows from a NoSQL table that satisfy a filter condition.

- Java
- Python
- Go
- Node.js
- C#

Java

You can use the DELETE SQL command in the Query request to delete data. To execute your query, you use the `NoSQLHandle.query()` API.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***ModifyData.java*** from the examples here.

```
//delete rows based on a filter condition
private static void deleteRows(NoSQLHandle handle, String sqlstmt) throws
Exception {
   QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
   handle.query(queryRequest);
   System.out.println("Deleted row(s) from table " + tableName);
}
```

```
String del_stmt ="DELETE FROM stream_acct acct1 WHERE
acct1.acct_data.firstName=\"Adelaide\" AND
acct1.acct_data.lastName=\"Willard\"";
/*delete rows based on a filter condition*/
deleteRows(handle,del_stmt);
```

Python

You can use the DELETE SQL command in the Query request to delete data. To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***ModifyData.py*** from the examples here.

```
#del row(s) with a filter condition
def delete_rows(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   result = handle.query(request)
   print('Deleted data from table: stream_acct')
```

```
# delete data based on a filter condition
del_stmt ='''DELETE FROM stream_acct acct1 WHERE
acct1.acct_data.firstName="Adelaide" AND acct1.acct_data.lastName="Willard"'''
delete_rows(handle,del_stmt)
```

Go

You can use the DELETE SQL command in the Query request to delete data. To execute a query use the `Client.Query` function.

Download the full code ***ModifyData.go*** from the examples here.

```
//delete rows based on a filter condition
func deleteRows(client *nosqldb.Client, err error, tableName string,
querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
      return
   }
   queryReq := &nosqldb.QueryRequest{
    PreparedStatement: &prepRes.PreparedStatement,    }
   var results []*types.MapValue
   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
      fmt.Printf("Upsert failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()
      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
      }
      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
   }
   for i, r := range results {
    fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
   }
   fmt.Printf("Deleted data from the table: %v\n",tableName)
}


delete_stmt := `DELETE FROM stream_acct acct1 WHERE
acct1.acct_data.firstName="Adelaide" AND acct1.acct_data.lastName="Willard"`
deleteRows(client, err,tableName,delete_stmt)
```

Node.js

You can use the DELETE SQL command in the Query request to delete data. To execute a query use `query` method.

JavaScript: Download the full code ***ModifyData.js*** from the examples here.

```
/*deletes data based on a filter conditioin */
async function deleteRows(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


await deleteRows(handle,del_stmt);
console.log("Rows deleted");
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

TypeScript: Download the full code ***ModifyData.ts*** from the examples here.

```
async function deleteRows(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


await deleteRows(handle,del_stmt);
console.log("Rows deleted");
```

C#

You can use the DELETE SQL command in the Query request to delete data. To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***ModifyData.cs*** from the examples here.

```
private static async Task deleteRows(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
}


await deleteRows(client,del_stmt);
Console.WriteLine("Rows removed from the table");
```

Simple SELECT queries

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Using Get API to fetch data
- Substituting column names in a query
- Using Query API to fetch data

Copyright © 2022, 2025, Oracle and/or its affiliates.

Using SQL commands to fetch data

You can use SQL SELECT statement to fetch data from your NoSQL table.

Fetching all rows from a table

You can choose columns from a table. To do so, list the names of the desired table columns after SELECT in the statement. You give the name of the table after the FROM clause. To retrieve data from a child table, use dot notation, such as **parent.child**. To choose all table columns, use the asterisk (*) wildcard character. The SELECT statement can also contain computational expressions based on the values of existing columns.

Example 1: Choose all data from the table `BaggageInfo`.

```
SELECT * FROM BaggageInfo
```

Explanation: The `BaggageInfo` schema has some fixed static fields and a JSON column. The static fields are ticket number, full name, gender, contact phone, and confirmation number. The bag information is stored as JSON and is populated with an array of documents.

Output (displaying only a row of the result for brevity):

```
{"ticketNo":1762330498104,"fullName":"Michelle
Payne","gender":"F","contactPhone":"575-781-6240","confNo":"RL3J4Q",
"bagInfo":[{
   "bagArrivalDate":"2019-02-02T23:59:00Z",
   "flightLegs":[
      {"actions":[
           {"actionAt":"SFO","actionCode":"ONLOAD to
IST","actionTime":"2019-02-02T12:10:00Z"},
           {"actionAt":"SFO","actionCode":"BagTag Scan at
SFO","actionTime":"2019-02-02T11:47:00Z"},
           {"actionAt":"SFO","actionCode":"Checkin at
SFO","actionTime":"2019-02-02T10:01:00Z"}],
           "estimatedArrival":"2019-02-03T01:00:00Z",
           "flightDate":"2019-02-02T12:00:00Z",
           "flightNo":"BM318",
           "fltRouteDest":"IST",
           "fltRouteSrc":"SFO"},
      {"actions":[
           {"actionAt":"IST","actionCode":"ONLOAD to
ATH","actionTime":"2019-02-03T13:06:00Z"},
           {"actionAt":"IST","actionCode":"BagTag Scan at
IST","actionTime":"2019-02-03T12:48:00Z"},
           {"actionAt":"IST","actionCode":"OFFLOAD from
IST","actionTime":"2019-02-03T13:00:00Z"}],
           "estimatedArrival":"2019-02-03T12:12:00Z",
           "flightDate":"2019-02-02T13:10:00Z",
           "flightNo":"BM696",
           "fltRouteDest":"ATH",
           "fltRouteSrc":"IST"},
      {"actions":[
           {"actionAt":"JTR","actionCode":"Offload to Carousel at
JTR","actionTime":"2019-02-03T00:06:00Z"},
           {"actionAt":"ATH","actionCode":"ONLOAD to
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
JTR","actionTime":"2019-02-03T00:13:00Z"},
           {"actionAt":"ATH","actionCode":"OFFLOAD from
ATH","actionTime":"2019-02-03T00:10:00Z"}],
           "estimatedArrival":"2019-02-03T00:12:00Z",
           "flightDate":"2019-2-2T12:10:00Z",
           "flightNo":"BM665",
           "fltRouteDest":"JTR",
           "fltRouteSrc":"ATH"}],
     "id":"79039899186259",
     "lastActionCode":"OFFLOAD",
     "lastActionDesc":"OFFLOAD",
     "lastSeenStation":"JTR",
     "lastSeenTimeGmt":"2019-02-02T23:59:00Z",
     "routing":"SFO/IST/ATH/JTR",
     "tagNum":"17657806247861"}
]}
```

Example 2: To choose specific column(s) from the table `BaggageInfo`, include the column names as a comma-separated list in the SELECT statement.

```
SELECT fullName, contactPhone, gender FROM BaggageInfo
```

Explanation: You want to display the values of three static fields - full name, phone number, and gender.

Output:

```
{"fullName":"Lucinda Beckman","contactPhone":"364-610-4444","gender":"M"}
{"fullName":"Adelaide Willard","contactPhone":"421-272-8082","gender":"M"}
{"fullName":"Raymond Griffin","contactPhone":"567-710-9972","gender":"F"}
{"fullName":"Elane Lemons","contactPhone":"600-918-8404","gender":"F"}
{"fullName":"Zina Christenson","contactPhone":"987-210-3029","gender":"M"}
{"fullName":"Zulema Martindale","contactPhone":"666-302-0028","gender":"F"}
{"fullName":"Dierdre Amador","contactPhone":"165-742-5715","gender":"M"}
{"fullName":"Henry Jenkins","contactPhone":"960-428-3843","gender":"F"}
{"fullName":"Rosalia Triplett","contactPhone":"368-769-5636","gender":"F"}
{"fullName":"Lorenzo Phil","contactPhone":"364-610-4444","gender":"M"}
{"fullName":"Gerard Greene","contactPhone":"395-837-3772","gender":"M"}
{"fullName":"Adam Phillips","contactPhone":"893-324-1064","gender":"M"}
{"fullName":"Doris Martin","contactPhone":"289-564-3497","gender":"F"}
{"fullName":"Joanne Diaz","contactPhone":"334-679-5105","gender":"F"}
{"fullName":"Omar Harvey","contactPhone":"978-191-8550","gender":"F"}
{"fullName":"Fallon Clements","contactPhone":"849-731-1334","gender":"M"}
{"fullName":"Lisbeth Wampler","contactPhone":"796-709-9501","gender":"M"}
{"fullName":"Teena Colley","contactPhone":"539-097-5220","gender":"M"}
{"fullName":"Michelle Payne","contactPhone":"575-781-6240","gender":"F"}
{"fullName":"Mary Watson","contactPhone":"131-183-0560","gender":"F"}
{"fullName":"Kendal Biddle","contactPhone":"619-956-8760","gender":"F"}
```

Example 3: Choose all data from the table `stream_acct`.

```
SELECT * FROM stream_acct
```

Explanation: The `stream_acct` schema has some fixed static fields and a JSON column.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output (displaying only a row of the result for brevity):

```
{"acct_id":1,"profile_name":"AP","account_expiry":"2023-10-18T00:00:00.0Z",
"acct_data":{
[{
   "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"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
         },
         {
            "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"
         }]
      }]
   }]}}
```

Filter data from a table

You can filter query results by specifying a filter condition in the WHERE clause. Typically, a filter condition consists of one or more comparison expressions connected through logical operators AND or OR. The following comparison operators are also supported: =, !=, >, >=, <, and <= .

Example 1: Find the tag number of a passenger's baggage along with the passenger's full name for a given reservation number **FH7G1W**.

```
SELECT bag.fullName, bag.bagInfo[].tagNum FROM BaggageInfo bag
WHERE bag.confNo="FH7G1W"
```

Explanation: You fetch the tag number corresponding to a given reservation number.

Output:

```
{"fullName":"Rosalia Triplett","tagNum":"17657806215913"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Note

For better understanding, the row of data with all the static fields and the bagInfo JSON is shown below.

```
"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"
      } ]
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        } ],
        "lastSeenTimeGmt" : "2019-02-02T03:13:00",
        "bagArrivalDate" : "2019.02.02T03:13:00"
      } ]
```

Example 2: Where was the baggage with a given reservation number **FH7G1W** last seen? Also, fetch the tag number of the baggage.

```
SELECT bag.fullName, bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenStation
FROM BaggageInfo bag WHERE bag.confNo="FH7G1W"
```

Explanation: The `bagInfo` is JSON and is populated with an array of documents. The full name and the last seen station can be fetched for a particular reservation number.

Output:

```
{"fullName":"Rosalia Triplett","tagNum":"17657806215913",
"lastSeenStation":"VIE"}
```

Example 3: Select details of the bags(tag and last seen time) for a passenger with ticket number **1762340579411**.

```
SELECT bag.ticketNo, bag.fullName,
bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenStation
FROM BaggageInfo bag where bag.ticketNo=1762320369957
```

Explanation: The `bagInfo` is JSON and is populated with an array of documents. The full name, tag number, and last seen station can be fetched for a particular ticket number.

Output:

```
{"fullName":"Lorenzo Phil","tagNum":["17657806240001","17657806340001"],
"lastSeenStation":["JTR","JTR"]}
```

Example 4: Fetch the last name, account expiry date and the shows watched by the user with `acct_id` 1.

```
SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName FROM stream_acct acct WHERE
acct_id=1
```

Explanation: The `acct_data` is JSON and is populated with an array of documents. The last name, account expiry date and show names are fetched for a particular account id.

Output:

```
{"account_expiry":"2023-10-18T00:00:00.0Z","lastName":"Phillips","showName":
["At the Ranch","Bienvenu"]}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Substituting column names in a query

You can use a different name for a column during a SELECT statement. Substituting a name in a query does not change the column name, but uses the substitute in the data returned.

Example: The following query returns the phone number as `CONTACTAT` in the result.

```
SELECT contactPhone AS CONTACTAT FROM BaggageInfo
```

Explanation: Here you want to fetch the contact phone of the passengers and display it as *CONTACTAT*.

Output:

```
{"CONTACTAT":"960-428-3843"}
{"CONTACTAT":"368-769-5636"}
{"CONTACTAT":"364-610-4444"}
{"CONTACTAT":"395-837-3772"}
{"CONTACTAT":"893-324-1064"}
{"CONTACTAT":"289-564-3497"}
{"CONTACTAT":"334-679-5105"}
{"CONTACTAT":"978-191-8550"}
{"CONTACTAT":"849-731-1334"}
{"CONTACTAT":"796-709-9501"}
{"CONTACTAT":"539-097-5220"}
{"CONTACTAT":"575-781-6240"}
{"CONTACTAT":"131-183-0560"}
{"CONTACTAT":"619-956-8760"}
{"CONTACTAT":"364-610-4444"}
{"CONTACTAT":"421-272-8082"}
{"CONTACTAT":"567-710-9972"}
{"CONTACTAT":"600-918-8404"}
{"CONTACTAT":"987-210-3029"}
{"CONTACTAT":"666-302-0028"}
{"CONTACTAT":"165-742-5715"}
```

You can combine columns using the concatenation operator "||" as shown below.

Example: For all customers, fetch the last place where the bag was seen and the time when it was seen.

Approach 1: Use the concatenation operator and fetch column names and static text as output of the SELECT command.

```
SELECT "The bag was last seen at " ||
bag.bagInfo[].lastSeenStation || " on " ||
bag.bagInfo[].bagArrivalDate AS Bag_Details FROM BaggageInfo bag
```

Output:

```
{"Bag_Details":"The bag was last seen at BZN on 2019-03-15T10:13:00Z"}
{"Bag_Details":"The bag was last seen at MEL on 2019-02-04T10:08:00Z"}
{"Bag_Details":"The bag was last seen at MEL on 2019-02-25T20:15:00Z"}
{"Bag_Details":"The bag was last seen at MAD on 2019-03-07T13:51:00Z"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"Bag_Details":"The bag was last seen at FRA on 2019-03-02T13:18:00Z"}
{"Bag_Details":"The bag was last seen at VIE on 2019-02-12T07:04:00Z"}
{"Bag_Details":"The bag was last seen at JTRJTR on
2019-03-12T15:05:00Z2019-03-12T16:25:00Z"}
{"Bag_Details":"The bag was last seen at JTR on 2019-03-07T16:01:00Z"}
{"Bag_Details":"The bag was last seen at MEL on 2019-02-01T16:13:00Z"}
{"Bag_Details":"The bag was last seen at MXP on 2019-03-22T10:17:00Z"}
{"Bag_Details":"The bag was last seen at MEL on 2019-02-16T16:13:00Z"}
{"Bag_Details":"The bag was last seen at MIA on 2019-03-02T16:09:00Z"}
{"Bag_Details":"The bag was last seen at BZN on 2019-02-21T14:08:00Z"}
{"Bag_Details":"The bag was last seen at SGN on 2019-02-10T10:01:00Z"}
{"Bag_Details":"The bag was last seen at JTR on 2019-02-02T23:59:00Z"}
{"Bag_Details":"The bag was last seen at BLR on 2019-03-14T06:22:00Z"}
{"Bag_Details":"The bag was last seen at VIE on 2019-03-05T12:00:00Z"}
{"Bag_Details":"The bag was last seen at JTR on 2019-03-12T15:05:00Z"}
{"Bag_Details":"The bag was last seen at SEA on 2019-02-15T21:21:00Z"}
{"Bag_Details":"The bag was last seen at HKG on 2019-02-03T08:09:00Z"}
{"Bag_Details":"The bag was last seen at HKG on 2019-02-13T11:15:00Z"}
```

The result is cluttered if there is more than one bag per customer/reservation number as shown above.

Approach 2: You can overcome this issue by printing as the value of elements of the `bagInfo` array as shown below.

```
SELECT "The bag was last seen at " || [bag.bagInfo[].lastSeenStation] || " on
" ||
[bag.bagInfo[].bagArrivalDate] AS Bag_Details FROM BaggageInfo bag
```

> **Note**
>
> Column names and static text can also be concatenated using the "||" operator.

Explanation: You are concatenating a part of the document in the `bagInfo` JSON with various static text and displaying it as elements of an array.

Output:

```
{"Bag_Details":"The bag was last seen at [\"MIA\"] on
[\"2019-03-02T16:09:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"BZN\"] on
[\"2019-02-21T14:08:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"SGN\"] on
[\"2019-02-10T10:01:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"HKG\"] on
[\"2019-02-13T11:15:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"JTR\"] on
[\"2019-02-02T23:59:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"BLR\"] on
[\"2019-03-14T06:22:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"VIE\"] on
[\"2019-03-05T12:00:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"JTR\"] on
[\"2019-03-12T15:05:00Z\"]"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"Bag_Details":"The bag was last seen at [\"SEA\"] on
[\"2019-02-15T21:21:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"HKG\"] on
[\"2019-02-03T08:09:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"BZN\"] on
[\"2019-03-15T10:13:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"MEL\"] on
[\"2019-02-04T10:08:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"MEL\"] on
[\"2019-02-25T20:15:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"MAD\"] on
[\"2019-03-07T13:51:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"FRA\"] on
[\"2019-03-02T13:18:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"VIE\"] on
[\"2019-02-12T07:04:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"JTR\",\"JTR\"] on
[\"2019-03-12T15:05:00Z\",
\"2019-03-12T16:25:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"JTR\"] on
[\"2019-03-07T16:01:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"MEL\"] on
[\"2019-02-01T16:13:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"MXP\"] on
[\"2019-03-22T10:17:00Z\"]"}
{"Bag_Details":"The bag was last seen at [\"MEL\"] on
[\"2019-02-16T16:13:00Z\"]"}
```

Using Get API to fetch data

Use the `GetRequest` API to fetch a single row of data using the primary key.

By default, all read operations are eventually consistent. This type of read is less costly than those using absolute consistency. You can change the default Consistency for a NoSQLHandle instance by using the `setConsistency()` method in the `Consistency` class.

- Java
- Python
- Go
- Node.js
- C#

Java

The `GetRequest` class provides a simple and powerful way to read data, while queries can be used for more complex read requests. To read data from a table, specify the target table and target key using the `GetRequest` class and use `NoSQLHandle.get()` to execute your request. The result of the operation is available in `GetResult`. You can change the default Consistency for a NoSQLHandle instance by using the

Copyright © 2022, 2025, Oracle and/or its affiliates.

`NoSQLHandleConfig.setConsistency(oracle.nosql.driver.Consistency)` and `GetRequest.setConsistency()` methods.

Download the full code ***QueryData.java*** from the examples here.

```
//Fetch single row using get API
private static void getRow(NoSQLHandle handle,String colName,int Id) throws
Exception {
   MapValue key = new MapValue().put(colName, Id);
   GetRequest getRequest = new GetRequest().setKey(key)
                                        .setTableName(tableName);
   GetResult getRes = handle.get(getRequest);
   /* on success, GetResult.getValue() returns a non-null value */
   if (getRes.getValue() != null) {
      System.out.println("\t" +getRes.getValue().toString());
   } else {
      System.out.println("Get Failed");
   }
}
```

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the `setTableName` method. Download the full code ***TableJoins.java*** from the examples to understand how to fetch data from a parent-child table here.

Python

Use the `GetRequest` API to fetch a single row of data using the primary key. You can read single rows using the `borneo.NoSQLHandle.get()` method. This method allows you to retrieve a record based on its primary key value. The `borneo.GetRequest` class is used for simple get operations. It contains the primary key value for the target row and returns an instance of `borneo.GetResult`. You can change the default Consistency for a `borneo.NoSQLHandle` using `borneo.NoSQLHandleConfig.set_consistency()` before creating the handle. It can be changed for a single request using `borneo.GetRequest.set_consistency()`.

Download the full code ***QueryData.py*** from the examples here.

```
# Fetch single row using get API
def getRow(handle,colName,Id):
  request = GetRequest().set_table_name('stream_acct')
  request.set_key({colName: Id})
  print('Query results: ')
  result = handle.get(request)
  print('Query results are' + str(result.get_value()))
```

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the `set_table_name` method. Download the full code ***TableJoins.py*** from the examples to understand how to fetch data from a parent-child table here.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Go

Use the `GetRequest` API to fetch a single row of data using the primary key. You can read single rows using the `Client.Get` function. This function allows you to retrieve a record based on its primary key value. The `nosqldb.GetRequest` is used for simple get operations. It contains the primary key value for the target row and returns an instance of `nosqldb.GetResult`. If the get operation succeeds, a non-nil `GetResult.Version` is returned. You can change the default Consistency for a `nosqldb.RequestConfig` using `RequestConfig.Consistency` before creating the client. It can be changed for a single request using `GetRequest.Consistency` field.

Download the full code ***QueryData.go*** from the examples here.

```
//fetch data from the table
func getRow(client *nosqldb.Client, err error, tableName string, colName
string, Id int)(){
    key:=&types.MapValue{}
   key.Put(colName, Id)
   req:=&nosqldb.GetRequest{
         TableName: tableName,
         Key: key,
   }
   res, err:=client.Get(req)
   if err != nil {
      fmt.Printf("GetResults() failed: %v\n", err)
      return
   }
   if res.RowExists() {
      fmt.Printf("Got row: %v\n", res.ValueAsJSON())
   } else {
      fmt.Printf("The row does not exist.\n")
   }
}
```

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the `TableName` parameter. Download the full code ***TableJoins.go*** from the examples to understand how to fetch data from a parent-child table here.

Node.js

Use the `get` API to fetch a single row of data using the primary key. You can read a single row using the `get` method. This method allows you to retrieve a record based on its primary key value. For method details, see NoSQLClient class.

Set the consistency of the read operation using Consistency enumeration. You can set the default Consistency for read operations in the initial configuration that is used to create a NoSQLClient instance using the consistency property. You can also change it for a single read operation by setting the consistency property in the GetOpt argument of the get method.

Copyright © 2022, 2025, Oracle and/or its affiliates.

JavaScript: Download the full code ***QueryData.js*** from the examples here.

```
//fetches single row with get API
async function getRow(handle,idVal) {
   try {
      const result = await handle.get(TABLE_NAME,{acct_Id: idVal });
      console.log('Got row:  %O', result.row);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

TypeScript: You can also supply an optional *type* parameter for a row to `get` and other data-related methods. The *type* parameter allows the TypeScript compiler to provide *type* hints for the returned `GetResult` method, as well as type-check the passed key. While *type* checking, the compiler inspects if the primary key field is a part of the table schema and if the *type* of the primary key field is one of the allowed *types*, that is either string, numeric, date or boolean. Download the full code ***QueryData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/*fetches single row with get API*/
async function getRow(handle: NoSQLClient,idVal: Integer) {
   try {
      const result = await handle.get<StreamInt>(TABLE_NAME,{acct_Id:
idVal });
      console.log('Got row:  %O', result.row);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the `TABLE_NAME` parameter. Download the full JavaScript code ***TableJoins.js*** here and the full TypeScript code ***TableJoins.ts*** here to understand how to fetch data from a parent-child table.

C#

Use the GET API to fetch a single row of data using the primary key. You can read a single row using the `GetAsync` method. This method allows you to retrieve a row based on its primary key value. The field names should be the same as the table primary key column names. You may also pass options as `GetOptions`.

You can set consistency of a read operation using Consistency enumeration. The default consistency for read operations may be set as Consistency property of `NoSQLConfig`. You may

Copyright © 2022, 2025, Oracle and/or its affiliates.

also change the consistency for a single Get operation by using Consistency property of `GetOptions`. `GetAsync` method returns `Task<GetResult<RecordValue>>`. `GetResult` instance contains the returned Row, the row Version and other information. If the row with the provided primary key does not exist in the table, the values of both Row and Version properties will be null.

Download the full code ***QueryData.cs*** from the examples here.

```
private static async Task getRow(NoSQLClient client,String colName, int Id){
   var result = await client.GetAsync(TableName,
         new MapValue
         {
             [colName] =Id
         });
   if (result.Row != null){
      Console.WriteLine("Got row: {0}\n", result.Row.ToJsonString());
   }
   else {
      Console.WriteLine("Row with primaryKey {0} doesn't exist",colName);
   }
}
```

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the `TableName` parameter. Download the full code ***TableJoins.cs*** from the examples to understand how to fetch data from a parent-child table here.

Using Query API to fetch data

You can use the `QueryRequest` to construct queries to filter data from your NoSQL table.

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API. See Oracle NoSQL Java SDK API Reference for more details on the various classes and methods.

There are two ways to get the results of a query: using an iterator or loop through partial results.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- **Iterator:** Use `NoSQLHandle.queryIterable(QueryRequest)` to get an iterable that contains all the results.
- **Partial Results:** To compute and retrieve the full result set of a query, the same `QueryRequest` instance will, in general, have to be executed multiple times (via `NoSQLHandle.query(oracle.nosql.driver.ops.QueryRequest)`. Each execution returns a `QueryRequest`, which contains a subset of the result set.

```
String sqlstmt_allrows="SELECT * FROM stream_acct";
private static void fetchRows(NoSQLHandle handle,String sqlstmt)
                                              throws Exception {
   try (
      QueryRequest queryRequest =
            new QueryRequest().setStatement(sqlstmt_allrows);

      QueryIterableResult results =
            handle.queryIterable(queryRequest)){

      for (MapValue res : results) {
         System.out.println("\t" + res);
      }
   }
}
```

You can also apply filter conditions using the WHERE clause in the query.

```
String sqlstmt_allrows=
"SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName
FROM stream_acct acct WHERE acct_id=1";
```

Download the full code **QueryData.java** from the examples here.

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the sql statement. Download the full code ***TableJoins.java*** from the examples to understand how to fetch data from a parent-child table here.

Python

To execute a query use the `borneo.NoSQLHandle.query()` method.

There are two ways to get the results of a query: using an iterator or loop through partial results.

- Use `borneo.NoSQLHandle.query_iterable()` to get an iterable that contains all the results of a query.
- You can loop through partial results by using the `borneo.NoSQLHandle.query()` method. For example, to execute a SELECT query to read data from your table, a `borneo.QueryResult` contains a list of results. And if the `borneo.QueryRequest.is_done()` returns False, there may be more results, so queries should generally be run in a loop. It is

Copyright © 2022, 2025, Oracle and/or its affiliates.

possible for single request to return no results but the query still not done, indicating that the query loop should continue.

```
sqlstmt = 'SELECT * FROM stream_acct'

def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)

   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))
```

You can also apply filter conditions using the WHERE clause in the query.

```
sqlstmt = 'SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName
FROM stream_acct acct WHERE acct_id=1'
```

Download the full code ***QueryData.py*** from the examples here.

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the sql statement. Download the full code ***TableJoins.py*** from the examples to understand how to fetch data from a parent-child table here.

Go

To execute a query use the `Client.Query` function. When execute on the cloud service, the amount of data read by a single query request is limited by a system default and can be further limited using `QueryRequest.MaxReadKB`. This limits the amount of data read and not the amount of data returned, which means that a query can return zero results but still have more data to read. For this reason queries should always operate in a loop, acquiring more results, until `QueryRequest.IsDone()` returns true, indicating that the query is done.

```
querystmt := "select * FROM stream_acct"

func fetchData(client *nosqldb.Client, err error,
               tableName string, querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{ Statement: querystmt,}

   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }

   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,}
   var results []*types.MapValue
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    for {
       queryRes, err := client.Query(queryReq)
       if err != nil {
          fmt.Printf("Query failed: %v\n", err)
       return
       }
       res, err := queryRes.GetResults()

       if err != nil {
          fmt.Printf("GetResults() failed: %v\n", err)
       return
       }

       results = append(results, res...)
       if queryReq.IsDone() {
          break
       }
    }
    for i, r := range results {
       fmt.Printf("\t%d: %s\n", i+1,
                    jsonutil.AsJSON(r.Map()))
    }
}
```

You can also apply filter conditions using the WHERE clause in the query.

```
querystmt := "SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName
FROM stream_acct acct where acct_id=1"
```

Download the full code ***QueryData.go*** from the examples here.

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the sql statement. Download the full code ***TableJoins.go*** from the examples to understand how to fetch data from a parent-child table here.

Node.js

You can query data from the NoSQL tables using one of these methods. For method details, see NoSQLClient class.

1. Use the `query` method to execute a query. This method returns a *Promise* of `QueryResult`,which is a plain JavaScript object containing an array of resulting rows as well as a continuation key. You can use the `query` method in two ways:
   - You can call the `query` method only once for queries that access at most one row. These queries can only include select statements based on the primary key (the where clause must specify equality based on the complete primary key). In all other cases, you can use either `query` in a loop or `queryIterable` method.
   - You can call the `query` method in a loop to retrieve multiple rows. As the amount of data returned by a query is limited by the system default and can be further limited by

Copyright © 2022, 2025, Oracle and/or its affiliates.

setting the `maxReadKB` property in the `QueryOpt` argument of the `query`one invocation of the `query` method can't return all the available results. To address this issue, run the query in a loop until the `continuationKey` in `QueryResult` becomes *null/undefined*.

2. Iterate over the query results using the `queryIterable` method. This method returns an iterable object that you can iterate over with a for-await-of loop. You need not manage the continuation in this method.

Note

With the `queryIterable`method, you can also use the `QueryOpt` argument with properties other than `continuationKey`.

JavaScript: Download the full code ***QueryData.js*** from the examples here.

```
const querystmt = 'SELECT * FROM stream_acct';

async function fetchData(handle,querystmt) {
    const opt = {};

    try {
        do {
            const result = await handle.query(querystmt, opt);
            for(let row of result.rows) {
                console.log('  %O', row);
            }

            opt.continuationKey = result.continuationKey;
        } while(opt.continuationKey);
    } catch(error) {
        console.error('  Error: ' + error.message);
    }
}
```

You can also apply filter conditions using the WHERE clause in the query.

```
const querystmt =
'SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName
FROM stream_acct acct WHERE acct_id=1';
```

TypeScript: You can use the same methods described in JavaScript above for TypeScript. You can also supply an optional query result schema as the *type* parameter to the `query` method to provide *type* hints for the rows returned in the `QueryResult`. This need not be the same as table row schema (unless using *SELECT ** query) as the query can include projections, name aliases, aggregate values, and so forth. Download the full code ***QueryData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

You can also apply filter conditions using the WHERE clause in the query.

```
const querystmt =
'SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName
FROM stream_acct acct WHERE acct_id=1';
```

> **Note**
>
> To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the sql statement. Download the full JavaScript code ***TableJoins.js*** here and the full TypeScript code ***TableJoins.ts*** here to understand how to fetch data from a parent-child table.

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable. You may pass options to each of these methods as `QueryOptions`. `QueryAsync` method return `Task<QueryResult<RecordValue>>`. `QueryResult` contains query results as a list of RecordValue instances, as well as other information. When your query specifies a complete primary key, it is sufficient to call `QueryAsync` once. The amount of data returned by the query is limited by the system. It could also be further limited by setting `MaxReadKB` property of `QueryOptions`. This means that one invocation of `QueryAsync` may not return all available results. This situation is dealt with by using continuation key. Non-null continuation key in `QueryResult` means that more query results may be available. This means that queries should run in a loop, looping until the continuation key becomes null. See Oracle NoSQL Dotnet SDK API Reference for more details of all classes and methods.

```
private const string querystmt ="SELECT * FROM stream_acct";

private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   await DoQuery(queryEnumerable);
 }

//function to display result
private static async Task
   DoQuery(IAsyncEnumerable<QueryResult<RecordValue>> queryEnumerable){
   Console.WriteLine("  Query results:");

   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}
```

You can also apply filter conditions using the WHERE clause in the query.

```
private const string querystmt =
"SELECT account_expiry, acct.acct_data.lastName,
acct.acct_data.contentStreamed[].showName
FROM stream_acct acct WHERE acct_id=1";
```

Download the full code ***QueryData.cs*** from the examples here.

Note

To fetch data from a child table, specify the full name of the table (parent_tablename.child_tablename) in the sql statement. Download the full code ***TableJoins.cs*** from the examples to understand how to fetch data from a parent-child table here.

SELECT queries on JSON collection tables

You can use the SQL expressions to query data from the JSON collection tables. The SQL queries work similarly on tables based on a fixed schema.

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 to the attribute.

To follow along with the examples, create a JSON collection table for a shopping application and insert the sample data records as described in the Sample data to run queries section. A few sample rows from the table are as follows:


```
{"contactPhone":"1517113582","address":
{"city":"Houston","number":651,"state":"TX","street":"Tex
Ave","zip":95085},"cart":null,"firstName":"Dierdre","lastName":"Amador","order
s":[{"EstDelivery":"2023-11-01","item":"handbag","orderID":"201200",
"priceperunit":350},
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"EstDelivery":"2023-11-01","item":"Lego","orderID":"201201","priceperunit":55
00}]}
```

```
{"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 1: 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 logical operator AND 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"
}
```

Example 2: 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"Message":"Hi Sharon","Offer":"The prices on Wallets have
    dropped"}
```

Using Path expressions

Path expressions are used to navigate inside hierarchically structured data. Oracle NoSQL Database supports different complex data types like arrays and records. You will learn how to work with different complex data types using path expressions.

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Using Internal variables and aliases
- Working with Arrays
- Working with nested data type
- Finding the size of a complex data type

Using Internal variables and aliases

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

The table name in a query 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.

Example 1: Find the ticket number and passenger details for a given reservation code:

```
SELECT bagDet.ticketNo, bagDet.fullName, bagDet.contactPhone FROM BaggageInfo
bagDet
WHERE confNo="QB100J"
```

Explanation: In this query, you fetch the values of static fields like fullname, ticket number, and contact phone for a particular reservation code. You use a table alias for the `BaggageInfo` table.

Output:

```
{"ticketNo":1762390789239,"fullName":"Zina
      Christenson","contactPhone":"987-210-3029"}
```

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.

Example 2: Fetch the full name and tag number for all customer baggage shipped after 2019.

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


Copyright © 2022, 2025, Oracle and/or its affiliates.

Explanation: The bag arrival date value for every bag should be greater than the year 2019. Here the "$element" is bound to the context row (every baggage of the customer). The EXISTS operator checks whether the sequence returned by its input expression is empty or not. The sequence returned by the comparison operator ">=" is non-empty for all bags which arrived after 2019.

Output:

```
{"fullName":"Lucinda Beckman","ticketNo":1762320569757}
{"fullName":"Adelaide Willard","ticketNo":1762392135540}
{"fullName":"Raymond Griffin","ticketNo":1762399766476}
{"fullName":"Elane Lemons","ticketNo":1762324912391}
{"fullName":"Zina Christenson","ticketNo":1762390789239}
{"fullName":"Zulema Martindale","ticketNo":1762340579411}
{"fullName":"Dierdre Amador","ticketNo":1762376407826}
{"fullName":"Henry Jenkins","ticketNo":176234463813}
{"fullName":"Rosalia Triplett","ticketNo":1762311547917}
{"fullName":"Lorenzo Phil","ticketNo":1762320369957}
{"fullName":"Gerard Greene","ticketNo":1762341772625}
{"fullName":"Adam Phillips","ticketNo":1762344493810}
{"fullName":"Doris Martin","ticketNo":1762355527825}
{"fullName":"Joanne Diaz","ticketNo":1762383911861}
{"fullName":"Omar Harvey","ticketNo":1762348904343}
{"fullName":"Fallon Clements","ticketNo":1762350390409}
{"fullName":"Lisbeth Wampler","ticketNo":1762355854464}
{"fullName":"Teena Colley","ticketNo":1762357254392}
{"fullName":"Michelle Payne","ticketNo":1762330498104}
{"fullName":"Mary Watson","ticketNo":1762340683564}
{"fullName":"Kendal Biddle","ticketNo":1762377974281}
```

Working with Arrays

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.

The `BaggageInfo` schema has many arrays. A simple array from the schema is the `actions` array in every `flightLeg`. You can use path expressions to navigate a simple array or a nested array.

```
"actions" : [ {
"actionAt" : "SYD",
"actionCode" : "ONLOAD to SIN",
"actionTime" : "2019.02.28 at 22:09:00 AEDT"
}, {
"actionAt" : "SYD",
"actionCode" : "BagTag Scan at SYD",
"actionTime" : "2019.02.28 at 21:51:00 AEDT"
}, {
"actionAt" : "SYD",
"actionCode" : "Checkin at SYD",
"actionTime" : "2019.02.28 at 20:06:00 AEDT"
} ]
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 1: Fetch the details of the first leg of every bag (including all the actions taken at the leg) for the passenger with ticket number **1762357254392**.

```
SELECT bagDet.fullName, bagDet.bagInfo[].flightLegs[0]
AS Details FROM BaggageInfo bagDet WHERE ticketNo=1762357254392
```

In the above query, `flightLegs` is an array. The slice step `[0]` is applied to the `flightLegs` array. Since array elements start with 0, this gives you the first record in the array. You get the first leg information of every bag for each passenger. You apply an additional filter with the `ticketNo` and so only one passenger information is fetched.

Output:

```
{"fullName":"Teena Colley",
"Details":[[
{"actionAt":"MSQ","actionCode":"ONLOAD to
FRA","actionTime":"2019-02-13T07:17:00Z"},
{"actionAt":"MSQ","actionCode":"BagTag Scan at
MSQ","actionTime":"2019-02-13T06:52:00Z"},
{"actionAt":"MSQ","actionCode":"Checkin at
MSQ","actionTime":"2019-02-13T06:11:00Z"}],
"2019-02-13T09:00:00Z","2019-02-13T07:00:00Z","BM365","FRA","MSQ"]}
```

> **Note**
>
> You can also use a slice step to select all array elements whose positions are within a range: [low: high], where low and high are expressions to specify the range boundaries. You can omit low and high expressions if you do not require a low or high boundary.

Example: Fetch the details of all the legs (including all the actions taken at all the legs) for the passenger with ticket number **1762357254392**.

You'll be using the slice step to fetch the first 3 records of the `flightLegs` array.

```
SELECT bagDet.fullName, bagDet.bagInfo[].flightLegs[0:2] AS Details
FROM BaggageInfo bagDet WHERE ticketNo=1762357254392
```

Output:

```
{"fullName":"Teena Colley",
"Details":[
 [
   {"actionAt":"MSQ","actionCode":"ONLOAD to
FRA","actionTime":"2019-02-13T07:17:00Z"},
   {"actionAt":"MSQ","actionCode":"BagTag Scan at
MSQ","actionTime":"2019-02-13T06:52:00Z"},
   {"actionAt":"MSQ","actionCode":"Checkin at
MSQ","actionTime":"2019-02-13T06:11:00Z"}
 ],
 "2019-02-13T09:00:00Z","2019-02-13T07:00:00Z","BM365","FRA","MSQ",
 [
   {"actionAt":"HKG","actionCode":"Offload to Carousel at
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
HKG","actionTime":"2019-02-13T11:15:00Z"},
   {"actionAt":"FRA","actionCode":"ONLOAD to
HKG","actionTime":"2019-02-13T10:39:00Z"},
   {"actionAt":"FRA","actionCode":"OFFLOAD from
FRA","actionTime":"2019-02-13T10:37:00Z"}
 ],
 "2019-02-13T11:18:00Z","2019-02-13T07:17:00Z","BM313","HKG","FRA"
]}
```

Working with nested data type

Oracle NoSQL database supports nested data type. That means you can have one data type inside another data type. For example, records inside an array, an array inside an array, and so on. The sample `Baggageinfo` schema uses nested data type of an array of arrays.

Example 1: Fetch the various actions taken on the first leg for the passenger with the ticket number **1762330498104**.

```
SELECT bagDet.fullName, bagDet.bagInfo[].flightLegs[0].values().values() AS
Action
FROM BaggageInfo bagDet WHERE ticketNo=1762330498104
```

Explanation: In the above query, `flightLegs` is a nested data type. This in turn has an `actions` array, which is an array of records. The above query is executed in two steps.

1. `$bag.bagInfo[].flightLegs[0].values()` gives all the entries in the first record of the`flightLegs` array. This includes an `actions` array. You can iterate this (using `values()`) to get all the records of the `actions` array as shown below.
2. `$bag.bagInfo[].flightLegs[0].values().values()` gives all the records of the `actions` array.

Output:

```
{"fullName":"Michelle Payne",
"Action":["SFO","ONLOAD to IST","2019-02-02T12:10:00Z","SFO",
"BagTag Scan at SFO","2019-02-02T11:47:00Z","SFO",
"Checkin at SFO","2019-02-02T10:01:00Z"]}
```

Example 2: Display details of the last transit action update done on the first leg for the passenger with the ticket number **1762340683564**.

```
SELECT bagDet.fullName, (bagDet.bagInfo[].flightLegs[0].values())
[2].actionCode
AS lastTransit_Update FROM BaggageInfo bagDet WHERE ticketNo=1762340683564
```

Explanation: The above query is processed using the following steps:

1. `$bagDet.bagInfo[].flightLegs[0].values()` gives all the entries in the first record of the `flightLegs` array.
2. `bagInfo[].flightLegs[0].values())[2]` points to the third (which is the last) record of the `actions` array inside the first element of the `flightLegs` array.
3. There are multiple records in the actions array. `bagInfo[].flightLegs[0].values())[2].actionCode` fetches the value corresponding to the `actionCode` element.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"fullName":"Mary Watson","lastTransit_Update":"Checkin at YYZ"}
```

Note

In a later section you will learn to write the same query in a generic way without hardcoding the array index by using the size function. See Finding the size of a complex data type.

Finding the size of a complex data type

The size function can be used to return the size (number of fields/entries) of a complex data type.

Example 1: Find out how many flight legs/hops are there for a passenger with ticket number **1762320569757**.

```
SELECT bagDet.fullName, size(bagDet.bagInfo.flightLegs) as Noof_Legs
FROM BaggageInfo bagDet WHERE ticketNo=1762320569757
```

Explanation: In the above query, you get the size of the `flightLegs` array using the `size` function.

Output:

```
{"fullName":"Lucinda Beckman","Noof_Legs":3}
```

Example 2: Find the number of action entries (for the bags) in the first leg for the passenger with ticket number **1762357254392**.

```
SELECT bagDet.fullName, size(bagDet.bagInfo[].flightLegs[0].actions) AS
FirstLeg_NoofActions
FROM BaggageInfo bagDet WHERE ticketNo=1762357254392
```

Output:

```
{"fullName":"Teena Colley","FirstLeg_NoofActions":3}
```

Example 3: Display details of the last transit action update done on the first leg for the passenger with the ticket number **1762340683564**.

```
SELECT bagDet.fullName,
(bagDet.bagInfo[].flightLegs[0].values())
[size(bagDet.bagInfo.flightLegs[0].actions)-1].actionCode
AS lastTransit_Update FROM BaggageInfo bagDet WHERE ticketNo=1762340683564
```

Output:

```
{"fullName":"Mary Watson","lastTransit_Update":"Checkin at YYZ"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Explanation:

The above query is processed using the following steps:

- 1. `$bagDet.bagInfo[].flightLegs[0].values()` gives all the entries in the first record of the `flightLegs` array.
  2. `size(bagDet.bagInfo.flightLegs[0].actions)`gives the size of the actions array in the first leg.
  3. There are multiple records in the actions array. You can use the result of the size function to get the last record in the action array and the corresponding `actionCode` can be fetched. You subtract the size by 1 as the index of an array starts with 0.

> **Note**
>
> The same query has been written in the topic Working with nested data type by hard coding the index of the actions array. Using the `size` function, you have rewritten the same query in a generic way without hard coding the index.

Using user-defined row metadata

> **Note**
>
> This feature has been made available to you on a "preview" basis so that you can get early access and provide feedback. It is intended for demonstration and preliminary use only. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to this feature and will not be responsible for any loss, costs, or damages incurred due to the use of this feature.

Oracle NoSQL Database supports the storing and retrieval of user-defined metadata for a row, allowing applications to associate custom metadata alongside the row's primary data. Although stored separately, this metadata is managed together with the row's primary data in an atomic manner. The database stores only the most recent state of each row and its associated metadata. When a row is updated, the change stream reflects the latest values for both the row and its metadata.

In addition to the system-managed metadata such as row's last modification time or TTL, you can use user-defined row metadata to add contextual information like audit tags, data clarifications, and so on without altering the core data model. This feature gives you more flexibility to annotate, track, analyze and extend your data, enabling richer functionality.

Applications supply this metadata as a JSON string, which may be a JSON object, array, string, number, boolean, or JSON null, and it is transactionally written along with the row. You can access this metadata through standard APIs or SQL functions.

Change Streams supports user-defined row metadata, allowing you to capture not only the changes made to the row's primary data but also any metadata that was supplied along with those changes. Streaming metadata enables applications to react intelligently and make context-aware decisions based on those changes.

- Using row metadata in Write Operations
- Using row metadata in Read Operations
- Using SQL commands on row metadata

Copyright © 2022, 2025, Oracle and/or its affiliates.

Using row metadata in Write Operations

You can supply your own metadata when you perform write operations on the rows in your tables with the help of the APIs below. For more information, see Using user-defined row metadata

This row metadata will be included in the change stream for understanding the context of operation. For more information, see Row metadata Streaming.

> **Note**
>
> This feature has been made available to you on a "preview" basis so that you can get early access and provide feedback. It is intended for demonstration and preliminary use only. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to this feature and will not be responsible for any loss, costs, or damages incurred due to the use of this feature.

- **Using PutRequest API**: You can use the `PutRequest` to add additional information, known as row metadata, alongside the main row data in your table. This row metadata is useful for storing extra information, that can describe or provide context about the actual row contents, such as the user who made the change or the source of the update. You must pass the row metadata as a JSON string using the `setRowMetadata()` method in the `PutRequest` class.
  This metadata will be included in the change stream event, allowing the change stream subscriber to see it.

  To understand this with an example, refer to `writeRowWithMetadata()` in ***ManageMetadata.java*** for more details, avaialable in the examples here.

```
/* Row Metadata inputs */
final static String rm1="{\"modified_by\" : \"John Doe\",\"reviewed_in\" :
\"Q1\",\"update_reason\" : \"Account details updated\"}";

/*Calling the method from the main function*/
writeRowWithMetadata(handle, (MapValue)newvalue, rm1);

/* Method to add row with associated row metadata*/
private static void writeRowWithMetadata(NoSQLHandle handle, MapValue
value, String rowMetadata) throws Exception {
   PutRequest putRequest = new PutRequest()
       .setValue(value)
       .setTableName(tableName)
       .setRowMetadata(rowMetadata);

   PutResult putResult = handle.put(putRequest);
   if (putResult.getVersion() != null) {
       System.out.println("Wrote: " + value);
   } else {
       System.out.println("Put failed");
      }
   }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

- **QueryRequest API**: When executing SQL DML statements using the QueryRequest API, you can apply `setRowMetadata()` to attach user-defined metadata to the affected rows. See the `updateRowViaQuery()` in ***ManageMetadata.java*** for more details, available in the examples here.
- **DeleteRequest API**: When deleting a single row using its primary key, you can attach metadata using `setRowMetadata()` to help describe the reason for deletion. See the `deleteRowWithMetadata()` in ***ManageMetadata.java*** for more details, available in the examples here.
- **MultiWrite API**: When performing atomic writes of multiple rows that share the same shard key, you can attach `setRowMetadata()` individually to each `PutRequest` in the batch. This is ideal for batch inserts or updates where each row may originate from a different user, source system, or import operation. See the `writeMultipleRows()` in ***MutiMetadataOps.java*** for more details, available in the examples here.
- **MultiDelete API**: When using a single atomic operation to delete multiple rows, that share the same shard key, you can attach the metadata using the `setRowMetadata()` in `MultiDeleteRequest` class. For example, you can use the metadata to document the reason for deleting those rows. See the `delMulRows()` in ***MutiMetadataOps.java*** for more details, available in the examples here.

Using row metadata in Read Operations

You can retrieve the user-defined row metadata with the help of APIs below. For more information, see Using user-defined row metadata

> **Note**
>
> This feature has been made available to you on a "preview" basis so that you can get early access and provide feedback. It is intended for demonstration and preliminary use only. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to this feature and will not be responsible for any loss, costs, or damages incurred due to the use of this feature.

- **GetRequest API**: You can use the `GetRequest` API to retrieve the row metadata of a single row. After executing the `GetRequest` API to retrieve the row, given a table name and primary key, you can retrieve the row metadata by calling the `getRowMetadata()` method. The operation returns a `GetResult` object that contains the retrieved data. If no metadata exists, it returns null.
  See the `readRowMetadata()` in ***ManageMetadata.java*** for more details, available in the examples here.

```
/*Method to fetch a row-metadata using GET API*/
private static void readRowMetadata(NoSQLHandle handle, int acctId) throws
Exception {
        GetRequest gr = new GetRequest()
            .setKey(new MapValue().put("acct_Id", acctId)) //key of the
row to fetch
            .setTableName(tableName);                     //table name

        GetResult gres = handle.get(gr);
        String rm= gres.getRowMetadata();

        if (rm != null) {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                System.out.println("Row metadata for acct_Id " + acctId + ": " 
+ rm);
            } else {
                System.out.println("No metadata found for acct_Id " + acctId);
            }
        }
```

- **QueryRequest API**: You can write a SQL SELECT statement using the QueryRequest API to explicitly fetch row metadata. This allows metadata to be returned as part of the result set. It's especially helpful when you want to view metadata across multiple rows for data clarification or analysis purposes. See the readAllRowMetadataViaQuery() in ***ManageMetadata.java*** for more details, available in the examples here.

Using SQL commands on row metadata

> **Note**
>
> This feature has been made available to you on a "preview" basis so that you can get early access and provide feedback. It is intended for demonstration and preliminary use only. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to this feature and will not be responsible for any loss, costs, or damages incurred due to the use of this feature.

With the SQL commands you can retrieve row metadata through queries, and you can also use row metadata in the WHERE clause to filter rows.

> **Note**
>
> You can only write row metadata through APIs, see Using user-defined row metadata.

Example: Fetch row metadata for all the rows in the table.

```
SELECT row_metadata($t) AS RowMetadata FROM stream_acct $t
```

Explanation: The query above fetches the metadata for all the rows in the stream_acct table.

Output:

```
{"RowMetadata":{"modified_by":"John 
Doe","reviewed_in":"Q1","update_reason":"Account details updated"}}
```

Example: Fetch the episode info of a user account based on the modified_by field in the row metadata

```
SELECT $t.acct_data.contentStreamed[1].seriesInfo[0].episodes[1] AS content 
FROM stream_acct $t WHERE row_metadata($t).modified_by="John Doe"
```

Exaplanation: In the query above, you fetch the episode details of a user account and use the row_metadata() to filter records that John Doe modified.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"content":
{"date":"2022-03-08","episodeID":30,"episodeName":"Merci","lengthMin":42,"minW
atched":42}}
```

Example: Delete data from the table using row metadata in the `WHERE` clause.

```
DELETE FROM stream_acct $t WHERE row_metadata($t).reviewed_in="Q1"
```

Explanation: The query above deletes the row from `stream_acct` table where `reviewed_in` field in the row's metadata has the value Q1.

Using Left Outer joins with parent-child tables

A JOIN is used to combine rows from two or more tables, based on a related column between them. In a hierarchical table, the 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. All tables in the hierarchy have the same shard key columns.

A Left Outer Join (LOJ) is one of the join operations that allows you to specify a join clause.

- Overview of Left Outer Joins
- Examples using Left Outer Joins

Overview of Left Outer Joins

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.

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.

Characteristics of LEFT OUTER JOIN:

- Queries multiple tables in the same hierarchy
- It is an ANSI-SQL Standard
- It does not support sibling table joins

If you want to follow along with the examples, download the script parentchildtbls_loaddata.sql and run it as shown below. This script creates the table used in the example 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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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 the `load` command, run the script.

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

Examples using Left Outer Joins

Various tables used in the examples :

- **`ticket`**

  ```
  ticketNo LONG
  confNo STRING
  PRIMARY KEY(ticketNo)
  ```

- **`ticket.bagInfo`**

  ```
  id LONG
  tagNum LONG
  routing STRING
  lastActionCode STRING
  lastActionDesc STRING
  lastSeenStation STRING,
  lastSeenTimeGmt TIMESTAMP(4)
  bagArrivalDate TIMESTAMP(4)
  PRIMARY KEY(id)
  ```

- **`ticket.bagInfo.flightLegs`**

  ```
  flightNo STRING
  flightDate TIMESTAMP(4)
  fltRouteSrc STRING
  fltRouteDest STRING
  estimatedArrival TIMESTAMP(4),
  actions JSON
  PRIMARY KEY(flightNo)
  ```

- **`ticket.passengerInfo`**

  ```
  contactPhone STRING
  fullName STRING
  gender STRING
  PRIMARY KEY(contactPhone)
  ```

- SQL Examples

Copyright © 2022, 2025, Oracle and/or its affiliates.

- Query API examples

SQL Examples

Example 1: Fetch the details of all passengers who have been issued a ticket.

```
SELECT fullname, contactPhone,gender
FROM ticket a
LEFT OUTER JOIN ticket.passengerInfo b
ON a.ticketNo=b.ticketNo
```

Explanation: This is an example of a join where the target table `ticket` is joined with its child table `passengerInfo`.

Output:

```
{"fullname":"Elane Lemons","contactPhone":"600-918-8404","gender":"F"}
{"fullname":"Adelaide Willard","contactPhone":"421-272-8082","gender":"M"}
{"fullname":"Dierdre Amador","contactPhone":"165-742-5715","gender":"M"}
{"fullname":"Doris Martin","contactPhone":"289-564-3497","gender":"F"}
{"fullname":"Adam Phillips","contactPhone":"893-324-1064","gender":"M"}
```

Example 1a: Fetch the details of the passenger with ticket number **1762324912391** .

```
SELECT fullname, contactPhone, gender
FROM ticket a
LEFT OUTER JOIN ticket.passengerInfo b
ON a.ticketNo=b.ticketNo
WHERE a.ticketNo=1762324912391
```

Explanation: This is an example of a join where the target table `ticket` is joined with its child table `passengerInfo` and a filter is applied to restrict the result. In this example, the result set is limited by applying a filter condition to the result of the join. You are limiting the result to a particular ticket number.

Output:

```
{"fullname":"Elane Lemons","contactPhone":"600-918-8404","gender":"F"}
```

Example 2: Fetch all the bag details for all passengers who have been issued a ticket.

```
SELECT * FROM ticket a
LEFT OUTER JOIN ticket.bagInfo b
ON a.ticketNo=b.ticketNo
```

Explanation: This is an example of a join where the target table`ticket` is joined with its child table `bagInfo`.

Output:

```
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":
{"ticketNo":1762344493810,"id":79039899165297,"tagNum":17657806255240,"routing
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
":"MIA/LAX/MEL",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MEL",
"lastSeenTimeGmt":"2019-02-01T16:13:00.0000Z","bagArrivalDate":"2019-02-01T16:
13:00.0000Z"}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
:"MXP/CDG/SLC/BZN",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"BZN",
"lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-15T10:
13:00.0000Z"}}

{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},
"b":
{"ticketNo":1762392135540,"id":79039899156435,"tagNum":17657806224224,"routing
":"GRU/ORD/SEA",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"SEA",
"lastSeenTimeGmt":"2019-02-15T21:21:00.0000Z","bagArrivalDate":"2019-02-15T21:
21:00.0000Z"}}

{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},
"b":
{"ticketNo":1762376407826,"id":7903989918469,"tagNum":17657806240229,"routing"
:"JFK/MAD",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MAD",
"lastSeenTimeGmt":"2019-03-07T13:51:00.0000Z","bagArrivalDate":"2019-03-07T13:
51:00.0000Z"}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"tagNum":17657806232501,"routing
":"BZN/SEA/CDG/MXP",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MXP",
"lastSeenTimeGmt":"2019-03-22T10:17:00.0000Z","bagArrivalDate":"2019-03-22T10:
17:00.0000Z"}}
```

Example 2a: Fetch all the bag details for a particular ticket number.

```
SELECT * FROM ticket a
LEFT OUTER JOIN ticket.bagInfo b
ON a.ticketNo=b.ticketNo
WHERE a.ticketNo=1762324912391
```

This is an example of a join where the target table `ticket` is joined with its child table `bagInfo` and a filter is applied to restrict the result. In this example, the result set is limited by applying a filter condition to the result of the join. You are limiting the result to a particular ticket number.

Output:

```
{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
:"MXP/CDG/SLC/BZN",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"BZN",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

"lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-15T10:13:00.0000Z"}}

> **Note**
>
> 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 ticket a
> LEFT OUTER JOIN ticket.bagInfo b
> ON a.ticketNo=b.ticketNo AND
> a.ticketNo=1762324912391
> ```

Output:


```
{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},"b":null}
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},"b":null}
{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"}, "b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
:"MXP/CDG/SLC/BZN",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"BZN",

"lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-15T10:
13:00.0000Z"}}
{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},"b":null}
{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},"b":null}
```


Example 3: Fetch all flight legs details for all passengers.

```
SELECT *FROM ticket a
LEFT OUTER JOIN ticket.bagInfo.flightLegs b
ON a.ticketNo=b.ticketNo;
```

Explanation: This is an example of a join where the target table `ticket` is joined with its descendant `ticketInfo`. A descendant table can be any level hierarchically below a table (For example `fightLegs` is the child of `bagInfo` which is the child of `ticket`, so `fightLegs` is a descendant of `ticket`).

Output:


```
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":
{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM604","flightDate":
"2019-02-01T06:00:00.0000Z",
"fltRouteSrc":"MIA","fltRouteDest":"LAX","estimatedArrival":"2019-02-01T11:00:
00.0000Z",
"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"},
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

{"actionAt":"MIA","actionCode":"Checkin at
MIA","actionTime":"2019-02-01T04:38:00Z"}]}}

{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":
{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM667","flightDate":
"2019-02-01T06:13:00.0000Z",
"fltRouteSrc":"LAX","fltRouteDest":"MEL","estimatedArrival":"2019-02-01T16:15:
00.0000Z",
"actions":[{"actionAt":"MEL","actionCode":"Offload to Carousel at
MEL","actionTime":"2019-02-01T16:15:00Z"},
{"actionAt":"LAX","actionCode":"ONLOAD to
MEL","actionTime":"2019-02-01T15:35:00Z"},
{"actionAt":"LAX","actionCode":"OFFLOAD from
LAX","actionTime":"2019-02-01T15:18:00Z"}]}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"flightNo":"BM170","flightDate":
"2019-03-15T08:13:00.0000Z",
"fltRouteSrc":"SLC","fltRouteDest":"BZN","estimatedArrival":"2019-03-15T10:14:
00.0000Z",
"actions":[{"actionAt":"BZN","actionCode":"Offload to Carousel at
BZN","actionTime":"2019-03-15T10:13:00Z"},
{"actionAt":"SLC","actionCode":"ONLOAD to
BZN","actionTime":"2019-03-15T10:06:00Z"},
{"actionAt":"SLC","actionCode":"OFFLOAD from
SLC","actionTime":"2019-03-15T09:59:00Z"}]}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"flightNo":"BM490","flightDate":
"2019-03-15T08:13:00.0000Z",
"fltRouteSrc":"CDG","fltRouteDest":"SLC","estimatedArrival":"2019-03-15T10:14:
00.0000Z",
"actions":[{"actionAt":"CDG","actionCode":"ONLOAD to
SLC","actionTime":"2019-03-15T09:42:00Z"},
{"actionAt":"CDG","actionCode":"BagTag Scan at
CDG","actionTime":"2019-03-15T09:17:00Z"},
{"actionAt":"CDG","actionCode":"OFFLOAD from
CDG","actionTime":"2019-03-15T09:19:00Z"}]}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"flightNo":"BM936","flightDate":
"2019-03-15T08:00:00.0000Z",
"fltRouteSrc":"MXP","fltRouteDest":"CDG","estimatedArrival":"2019-03-15T09:00:
00.0000Z",
"actions":[{"actionAt":"MXP","actionCode":"ONLOAD to
CDG","actionTime":"2019-03-15T08:13:00Z"},
{"actionAt":"MXP","actionCode":"BagTag Scan at
MXP","actionTime":"2019-03-15T07:48:00Z"},
{"actionAt":"MXP","actionCode":"Checkin at
MXP","actionTime":"2019-03-15T07:38:00Z"}]}}

{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},

Copyright © 2022, 2025, Oracle and/or its affiliates.

```
"b":
{"ticketNo":1762392135540,"id":79039899156435,"flightNo":"BM79","flightDate":"
2019-02-15T01:00:00.0000Z",
"fltRouteSrc":"GRU","fltRouteDest":"ORD","estimatedArrival":"2019-02-15T11:00:
00.0000Z",
"actions":[{"actionAt":"GRU","actionCode":"ONLOAD to
ORD","actionTime":"2019-02-15T01:21:00Z"},
{"actionAt":"GRU","actionCode":"BagTag Scan at
GRU","actionTime":"2019-02-15T00:55:00Z"},
{"actionAt":"GRU","actionCode":"Checkin at
GRU","actionTime":"2019-02-14T23:49:00Z"}]}}

{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"}
,"b":
{"ticketNo":1762392135540,"id":79039899156435,"flightNo":"BM907","flightDate":
"2019-02-15T01:21:00.0000Z",
"fltRouteSrc":"ORD","fltRouteDest":"SEA","estimatedArrival":"2019-02-15T21:22:
00.0000Z",
"actions":[{"actionAt":"SEA","actionCode":"Offload to Carousel at
SEA","actionTime":"2019-02-15T21:16:00Z"},
{"actionAt":"ORD","actionCode":"ONLOAD to
SEA","actionTime":"2019-02-15T20:52:00Z"},
{"actionAt":"ORD","actionCode":"OFFLOAD from
ORD","actionTime":"2019-02-15T20:44:00Z"}]}}

{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},
"b":
{"ticketNo":1762376407826,"id":7903989918469,"flightNo":"BM495","flightDate":"
2019-03-07T07:00:00.0000Z",
"fltRouteSrc":"JFK","fltRouteDest":"MAD","estimatedArrival":"2019-03-07T14:00:
00.0000Z",
"actions":[{"actionAt":"MAD","actionCode":"Offload to Carousel at
MAD","actionTime":"2019-03-07T13:54:00Z"},
{"actionAt":"JFK","actionCode":"ONLOAD to
MAD","actionTime":"2019-03-07T07:00:00Z"},
{"actionAt":"JFK","actionCode":"BagTag Scan at
JFK","actionTime":"2019-03-07T06:53:00Z"},
{"actionAt":"JFK","actionCode":"Checkin at
JFK","actionTime":"2019-03-07T05:03:00Z"}]}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"flightNo":"BM386","flightDate":
"2019-03-22T07:23:00.0000Z",
"fltRouteSrc":"CDG","fltRouteDest":"MXP","estimatedArrival":"2019-03-22T10:24:
00.0000Z",
"actions":[{"actionAt":"MXP","actionCode":"Offload to Carousel at
MXP","actionTime":"2019-03-22T10:15:00Z"},
{"actionAt":"CDG","actionCode":"ONLOAD to
MXP","actionTime":"2019-03-22T10:09:00Z"},
{"actionAt":"CDG","actionCode":"OFFLOAD from
CDG","actionTime":"2019-03-22T10:01:00Z"}]}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"flightNo":"BM578","flightDate":
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

"2019-03-22T07:23:00.0000Z",
"fltRouteSrc":"SEA","fltRouteDest":"CDG","estimatedArrival":"2019-03-21T23:24:00.0000Z",
"actions":[{"actionAt":"SEA","actionCode":"ONLOAD to CDG","actionTime":"2019-03-22T11:26:00Z"},
{"actionAt":"SEA","actionCode":"BagTag Scan at SEA","actionTime":"2019-03-22T10:57:00Z"},
{"actionAt":"SEA","actionCode":"OFFLOAD from SEA","actionTime":"2019-03-22T11:07:00Z"}]}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"flightNo":"BM704","flightDate":"2019-03-22T07:00:00.0000Z",
"fltRouteSrc":"BZN","fltRouteDest":"SEA","estimatedArrival":"2019-03-22T09:00:00.0000Z",
"actions":[{"actionAt":"BZN","actionCode":"ONLOAD to SEA","actionTime":"2019-03-22T07:23:00Z"},
{"actionAt":"BZN","actionCode":"BagTag Scan at BZN","actionTime":"2019-03-22T06:58:00Z"},
{"actionAt":"BZN","actionCode":"Checkin at BZN","actionTime":"2019-03-22T05:20:00Z"}]}}

Example 3a: Fetch all the flight leg details for a particular ticket number.

```
SELECT * FROM ticket a
LEFT OUTER JOIN ticket.bagInfo.flightLegs b
ON a.ticketNo=b.ticketNo
WHERE a.ticketNo=1762344493810
```

This is an example of a join where the target table `ticket` is joined with its descendant `bagInfo` and a filter is applied to restrict the result. In this example, the result set is limited by applying a filter condition to the result of the join. You are limiting the result to a particular ticket number.

The result has two rows, implying there are two flight legs for this ticket number.

Output:

"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM604",
"flightDate":"2019-02-01T06:00:00.0000Z","fltRouteSrc":"MIA","fltRouteDest":"LAX",
"estimatedArrival":"2019-02-01T11:00:00.0000Z",
"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"}]}}

{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM667",
"flightDate":"2019-02-01T06:13:00.0000Z","fltRouteSrc":"LAX","fltRouteDest":"MEL",
"estimatedArrival":"2019-02-01T16:15:00.0000Z",

Copyright © 2022, 2025, Oracle and/or its affiliates.

```
"actions":[{"actionAt":"MEL","actionCode":"Offload to Carousel at
MEL","actionTime":"2019-02-01T16:15:00Z"},
{"actionAt":"LAX","actionCode":"ONLOAD to
MEL","actionTime":"2019-02-01T15:35:00Z"},
{"actionAt":"LAX","actionCode":"OFFLOAD from
LAX","actionTime":"2019-02-01T15:18:00Z"}]}}
```

Example 4: Fetch the bag id and number of hops for all bags of all passengers.

```
SELECT b.id,count(*) AS NUMBER_HOPS
FROM ticket a LEFT OUTER JOIN ticket.bagInfo.flightLegs b
ON a.ticketNo=b.ticketNo GROUP BY b.id
```

Explanation: You group the data based on the bag id (using GROUP BY) and get the count of flight legs (using count()) for every bag.

Output:

```
{"id":79039899168383,"NUMBER_HOPS":3}
{"id":79039899156435,"NUMBER_HOPS":2}
{"id":7903989918469,"NUMBER_HOPS":1}
{"id":79039899165297,"NUMBER_HOPS":2}
{"id":79039899197492,"NUMBER_HOPS":3}
```

Example 4a: Find the number of hops for all the bags of a given passenger.

```
SELECT b.id,count(*) AS NUMBER_HOPS
FROM ticket a LEFT OUTER JOIN ticket.bagInfo.flightLegs b
ON a.ticketNo=b.ticketNo
WHERE a.ticketNo=1762355527825 GROUP BY b.id
```

Explanation: You group the data based on the bag id (using GROUP BY) and get the count of flight legs (Using count())for every bag. Additionally, you filter the results for a particular ticket number.

Output:

```
{"id":79039899197492,"NUMBER_HOPS":3}
```

Example 5: Fetch bag id and routing details of all bags that arrived after 2019.

```
SELECT b.id, routing
FROM ticket a LEFT OUTER JOIN ticket.bagInfo b
ON a.ticketNo=b.ticketNo
WHERE CAST (b.bagArrivalDate AS Timestamp(0))
>= CAST ("2019-01-01T00:00:00" AS Timestamp(0))
```

Explanation: This is an example of a join where the target table`ticket` is joined with its child table `bagInfo`. The filter condition is applied on the `bagArrivalDate`. The `CAST` function is used to convert the string into Timestamp and then the values are compared.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"id":79039899197492,"routing":"BZN/SEA/CDG/MXP"}
{"id":79039899165297,"routing":"MIA/LAX/MEL"}
{"id":79039899168383,"routing":"MXP/CDG/SLC/BZN"}
{"id":79039899156435,"routing":"GRU/ORD/SEA"}
{"id":7903989918469,"routing":"JFK/MAD"}
```

Query API examples

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***TableJoins.java*** from the examples here.

```
   /* fetch rows based on joins*/
private static void fetchRows(NoSQLHandle handle,String sql_stmt) throws
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sql_stmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)) {
         System.out.println("Query results:");
         for (MapValue res : results) {
            System.out.println("\t" + res);
         }
      }
}

/* fetching rows using left outer joins*/
String sql_stmt_loj ="SELECT * FROM ticket a LEFT OUTER JOIN
ticket.bagInfo.flightLegs b ON a.ticketNo=b.ticketNo";
System.out.println("Fetching data using Left outer joins:");
fetchRows(handle,sql_stmt_loj);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***TableJoins.py*** from the examples here.

```
# Fetch data from the table based on joins
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))

sql_stmt_loj='SELECT * FROM ticket a LEFT OUTER JOIN
ticket.bagInfo.flightLegs b ON a.ticketNo=b.ticketNo'
print('Fetching data using Left Outer Joins ')
fetch_data(handle,sql_stmt_loj)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***TableJoins.go*** from the examples here.

```
func fetchData(client *nosqldb.Client, err error,
               tableName string, querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{ Statement: querystmt,}

   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }

   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,}
   var results []*types.MapValue

   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()

      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
      }

      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
   }
   for i, r := range results {
      fmt.Printf("\t%d: %s\n", i+1,
                  jsonutil.AsJSON(r.Map()))
   }
}

querystmt_loj:= "SELECT * FROM ticket a LEFT OUTER JOIN
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
ticket.bagInfo.flightLegs b ON a.ticketNo=b.ticketNo"
fmt.Println("Fetching data using Left Outer Joins")
fetchData(client, err,querystmt_loj)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***TableJoins.js*** from the examples here.

```
//fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}

const stmt_loj = 'SELECT * FROM ticket a LEFT OUTER JOIN
ticket.bagInfo.flightLegs b ON a.ticketNo=b.ticketNo';
console.log("Fetching data using Left Outer Joins");
await fetchData(handle,stmt_loj);
```

TypeScript: Download the full code ***TableJoins.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
}

const stmt_loj = 'SELECT * FROM ticket a LEFT OUTER JOIN
ticket.bagInfo.flightLegs b ON a.ticketNo=b.ticketNo';
console.log("Fetching data using Left Outer Joins");
await fetchData(handle,stmt_loj);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***TableJoins.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Row
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}

private const string stmt_loj ="SELECT * FROM ticket a LEFT OUTER JOIN
ticket.bagInfo.flightLegs b ON a.ticketNo=b.ticketNo";
Console.WriteLine("Fetching data using Left Outer Joins: ");
await fetchData(client,stmt_loj);
```

Using NESTED TABLES to join parent-child tables

A JOIN is used to combine rows from two or more tables, based on a related column between them. In a hierarchical table, the 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. All tables in the hierarchy have the same shard key columns.

You can use NESTED TABLES clause to join tables in Oracle NoSQL Database.

- Overview of NESTED TABLES
- Examples using NESTED TABLES

Overview of NESTED TABLES

The NESTED TABLES clause specifies the participating tables and separates them into 3 groups. First, the target table from where the data is fetched is specified. Then the

Copyright © 2022, 2025, Oracle and/or its affiliates.

ANCESTORS clause, if present, specifies the number of tables that must be ancestors of the target table in the table hierarchy. Finally, the DESCENDANTS clause, if present, specifies the number of tables that must be descendants of the target table in the table hierarchy.

Note

Semantically, a NESTED TABLES clause is equivalent to a number of left-outer-join operations "centered" around the target table.

Characteristics of NESTED tables:

- Queries multiple tables in the same hierarchy
- It is not an ANSI-SQL Standard
- It supports sibling tables join

Table 4-1 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

If you want to follow along with the examples, download the script parentchildtbls_loaddata.sql and run it as shown below. This script creates the table used in the example 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
```

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 the `load` command, run the script.

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

Examples using NESTED TABLES

Various tables used in the examples :

Copyright © 2022, 2025, Oracle and/or its affiliates.

- **ticket**

```
ticketNo LONG
confNo STRING
PRIMARY KEY(ticketNo)
```

- **ticket.bagInfo**

```
id LONG
tagNum LONG
routing STRING
lastActionCode STRING
lastActionDesc STRING
lastSeenStation STRING,
lastSeenTimeGmt TIMESTAMP(4)
bagArrivalDate TIMESTAMP(4)
PRIMARY KEY(id)
```

- **ticket.bagInfo.flightLegs**

```
flightNo STRING
flightDate TIMESTAMP(4)
fltRouteSrc STRING
fltRouteDest STRING
estimatedArrival TIMESTAMP(4),
actions JSON
PRIMARY KEY(flightNo)
```

- **ticket.passengerInfo**

```
contactPhone STRING
fullName STRING
gender STRING
PRIMARY KEY(contactPhone)
```

- SQL Examples
- Query API examples

SQL Examples

Example 1: Fetch the details of all passengers who have been issued a ticket.

```
SELECT fullname, contactPhone, gender
FROM NESTED TABLES
(ticket a descendants(ticket.passengerInfo b))
```

Explanation: This is an example of a join where the target table `ticket` is joined with its child table `passengerInfo`.

Output:

```
{"fullname":"Elane Lemons","contactPhone":"600-918-8404","gender":"F"}
{"fullname":"Adelaide Willard","contactPhone":"421-272-8082","gender":"M"}
{"fullname":"Dierdre Amador","contactPhone":"165-742-5715","gender":"M"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"fullname":"Doris Martin","contactPhone":"289-564-3497","gender":"F"}
{"fullname":"Adam Phillips","contactPhone":"893-324-1064","gender":"M"}
```

Example 1a: Fetch the details of the passenger with ticket number **1762324912391** .

```
SELECT fullname, contactPhone, gender
FROM NESTED TABLES
(ticket a descendants(ticket.passengerInfo b))
WHERE a.ticketNo=1762324912391
```

Explanation: This is an example of a join where the target table`ticket` is joined with its child table `passengerInfo`. Additionally, you can limit the result set by applying a filter condition to the result of the join. You are limiting the result to a particular ticket number.

Output:

```
{"fullname":"Elane Lemons","contactPhone":"600-918-8404","gender":"F"}
```

Example 2: Fetch all the bag details for all passengers who have been issued a ticket.

```
SELECT * FROM NESTED TABLES
(ticket a descendants(ticket.bagInfo b))
```

Explanation: This is an example of a join where the target table `ticket` is joined with its child table `bagInfo`.

Output:

```
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":
{"ticketNo":1762344493810,"id":79039899165297,"tagNum":17657806255240,"routing
":"MIA/LAX/MEL",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MEL",
"lastSeenTimeGmt":"2019-02-01T16:13:00.0000Z","bagArrivalDate":"2019-02-01T16:
13:00.0000Z"}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
:"MXP/CDG/SLC/BZN",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"BZN",
"lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-15T10:
13:00.0000Z"}}

{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},
"b":
{"ticketNo":1762392135540,"id":79039899156435,"tagNum":17657806224224,"routing
":"GRU/ORD/SEA",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"SEA",
"lastSeenTimeGmt":"2019-02-15T21:21:00.0000Z","bagArrivalDate":"2019-02-15T21:
21:00.0000Z"}}

{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},
"b":
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"ticketNo":1762376407826,"id":7903989918469,"tagNum":17657806240229,"routing"
:"JFK/MAD",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MAD",
"lastSeenTimeGmt":"2019-03-07T13:51:00.0000Z","bagArrivalDate":"2019-03-07T13:
51:00.0000Z"}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"tagNum":17657806232501,"routing
":"BZN/SEA/CDG/MXP",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"MXP",
"lastSeenTimeGmt":"2019-03-22T10:17:00.0000Z","bagArrivalDate":"2019-03-22T10:
17:00.0000Z"}}
```

Example 2a: Fetch all the bag details for a particular ticket number.

```
SELECT * FROM
NESTED TABLES (ticket a descendants(ticket.bagInfo b))
WHERE a.ticketNo=1762324912391
```

Explanation: This is an example of a join where the target table`ticket` is joined with its child table `bagInfo`. Additionally, you can limit the result set by applying a filter condition to the result of the join. You are limiting the result to a particular ticket number.

Output:

```
{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
:"MXP/CDG/SLC/BZN",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"BZN",
"lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-15T10:
13:00.0000Z"}}
```

> **Note**
>
> 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
> NESTED TABLES(ticket a descendants(ticket.bagInfo b
> ON a.ticketNo=b.ticketNo
> AND a.ticketNo=1762324912391))
> ```

Output:

```
{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},"b":null}
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},"b":null}
{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"}, "b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
:"MXP/CDG/SLC/BZN",
"lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"BZN",

"lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-15T10:
13:00.0000Z"}}
{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},"b":null}
{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},"b":null}
```

Example 3: Fetch all flight leg details for all passengers.

```
SELECT * FROM
NESTED TABLES (ticket a descendants(ticket.bagInfo.flightLegs b))
```

Explanation: This is an example of a join where the target table `ticket` is joined with its descendant `bagInfo`. A descendant table can be any level hierarchically below a table (For example `fightLegs` is the child of `bagInfo` which is the child of `ticket`, so `flightLegs` is a descendant of `ticket`). All the rows from the `ticket` table will be fetched. If any row from the `ticket` table does not have a matching row in the `flightLegs` table, then NULL values will be displayed for those rows of the `flightLegs` table.

Output:

```
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":
{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM604","flightDate":
"2019-02-01T06:00:00.0000Z",
"fltRouteSrc":"MIA","fltRouteDest":"LAX","estimatedArrival":"2019-02-01T11:00:
00.0000Z",
"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"}]}}

{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":
{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM667","flightDate":
"2019-02-01T06:13:00.0000Z",
"fltRouteSrc":"LAX","fltRouteDest":"MEL","estimatedArrival":"2019-02-01T16:15:
00.0000Z",
"actions":[{"actionAt":"MEL","actionCode":"Offload to Carousel at
MEL","actionTime":"2019-02-01T16:15:00Z"},
{"actionAt":"LAX","actionCode":"ONLOAD to
MEL","actionTime":"2019-02-01T15:35:00Z"},
{"actionAt":"LAX","actionCode":"OFFLOAD from
LAX","actionTime":"2019-02-01T15:18:00Z"}]}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"flightNo":"BM170","flightDate":
"2019-03-15T08:13:00.0000Z",
"fltRouteSrc":"SLC","fltRouteDest":"BZN","estimatedArrival":"2019-03-15T10:14:
00.0000Z",
"actions":[{"actionAt":"BZN","actionCode":"Offload to Carousel at
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
BZN","actionTime":"2019-03-15T10:13:00Z"},
{"actionAt":"SLC","actionCode":"ONLOAD to
BZN","actionTime":"2019-03-15T10:06:00Z"},
{"actionAt":"SLC","actionCode":"OFFLOAD from
SLC","actionTime":"2019-03-15T09:59:00Z"}]}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"flightNo":"BM490","flightDate":
"2019-03-15T08:13:00.0000Z",
"fltRouteSrc":"CDG","fltRouteDest":"SLC","estimatedArrival":"2019-03-15T10:14:
00.0000Z",
"actions":[{"actionAt":"CDG","actionCode":"ONLOAD to
SLC","actionTime":"2019-03-15T09:42:00Z"},
{"actionAt":"CDG","actionCode":"BagTag Scan at
CDG","actionTime":"2019-03-15T09:17:00Z"},
{"actionAt":"CDG","actionCode":"OFFLOAD from
CDG","actionTime":"2019-03-15T09:19:00Z"}]}}

{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},
"b":
{"ticketNo":1762324912391,"id":79039899168383,"flightNo":"BM936","flightDate":
"2019-03-15T08:00:00.0000Z",
"fltRouteSrc":"MXP","fltRouteDest":"CDG","estimatedArrival":"2019-03-15T09:00:
00.0000Z",
"actions":[{"actionAt":"MXP","actionCode":"ONLOAD to
CDG","actionTime":"2019-03-15T08:13:00Z"},
{"actionAt":"MXP","actionCode":"BagTag Scan at
MXP","actionTime":"2019-03-15T07:48:00Z"},
{"actionAt":"MXP","actionCode":"Checkin at
MXP","actionTime":"2019-03-15T07:38:00Z"}]}}

{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},
"b":
{"ticketNo":1762392135540,"id":79039899156435,"flightNo":"BM79","flightDate":"
2019-02-15T01:00:00.0000Z",
"fltRouteSrc":"GRU","fltRouteDest":"ORD","estimatedArrival":"2019-02-15T11:00:
00.0000Z",
"actions":[{"actionAt":"GRU","actionCode":"ONLOAD to
ORD","actionTime":"2019-02-15T01:21:00Z"},
{"actionAt":"GRU","actionCode":"BagTag Scan at
GRU","actionTime":"2019-02-15T00:55:00Z"},
{"actionAt":"GRU","actionCode":"Checkin at
GRU","actionTime":"2019-02-14T23:49:00Z"}]}}

{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"}
,"b":
{"ticketNo":1762392135540,"id":79039899156435,"flightNo":"BM907","flightDate":
"2019-02-15T01:21:00.0000Z",
"fltRouteSrc":"ORD","fltRouteDest":"SEA","estimatedArrival":"2019-02-15T21:22:
00.0000Z",
"actions":[{"actionAt":"SEA","actionCode":"Offload to Carousel at
SEA","actionTime":"2019-02-15T21:16:00Z"},
{"actionAt":"ORD","actionCode":"ONLOAD to
SEA","actionTime":"2019-02-15T20:52:00Z"},
{"actionAt":"ORD","actionCode":"OFFLOAD from
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
ORD","actionTime":"2019-02-15T20:44:00Z"}]}}

{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},
"b":
{"ticketNo":1762376407826,"id":7903989918469,"flightNo":"BM495","flightDate":"
2019-03-07T07:00:00.0000Z",
"fltRouteSrc":"JFK","fltRouteDest":"MAD","estimatedArrival":"2019-03-07T14:00:
00.0000Z",
"actions":[{"actionAt":"MAD","actionCode":"Offload to Carousel at
MAD","actionTime":"2019-03-07T13:54:00Z"},
{"actionAt":"JFK","actionCode":"ONLOAD to
MAD","actionTime":"2019-03-07T07:00:00Z"},
{"actionAt":"JFK","actionCode":"BagTag Scan at
JFK","actionTime":"2019-03-07T06:53:00Z"},
{"actionAt":"JFK","actionCode":"Checkin at
JFK","actionTime":"2019-03-07T05:03:00Z"}]}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"flightNo":"BM386","flightDate":
"2019-03-22T07:23:00.0000Z",
"fltRouteSrc":"CDG","fltRouteDest":"MXP","estimatedArrival":"2019-03-22T10:24:
00.0000Z",
"actions":[{"actionAt":"MXP","actionCode":"Offload to Carousel at
MXP","actionTime":"2019-03-22T10:15:00Z"},
{"actionAt":"CDG","actionCode":"ONLOAD to
MXP","actionTime":"2019-03-22T10:09:00Z"},
{"actionAt":"CDG","actionCode":"OFFLOAD from
CDG","actionTime":"2019-03-22T10:01:00Z"}]}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"flightNo":"BM578","flightDate":
"2019-03-22T07:23:00.0000Z",
"fltRouteSrc":"SEA","fltRouteDest":"CDG","estimatedArrival":"2019-03-21T23:24:
00.0000Z",
"actions":[{"actionAt":"SEA","actionCode":"ONLOAD to
CDG","actionTime":"2019-03-22T11:26:00Z"},
{"actionAt":"SEA","actionCode":"BagTag Scan at
SEA","actionTime":"2019-03-22T10:57:00Z"},
{"actionAt":"SEA","actionCode":"OFFLOAD from
SEA","actionTime":"2019-03-22T11:07:00Z"}]}}

{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},
"b":
{"ticketNo":1762355527825,"id":79039899197492,"flightNo":"BM704","flightDate":
"2019-03-22T07:00:00.0000Z",
"fltRouteSrc":"BZN","fltRouteDest":"SEA","estimatedArrival":"2019-03-22T09:00:
00.0000Z",
"actions":[{"actionAt":"BZN","actionCode":"ONLOAD to
SEA","actionTime":"2019-03-22T07:23:00Z"},
{"actionAt":"BZN","actionCode":"BagTag Scan at
BZN","actionTime":"2019-03-22T06:58:00Z"},
{"actionAt":"BZN","actionCode":"Checkin at
BZN","actionTime":"2019-03-22T05:20:00Z"}]}}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 3a: Fetch all the flight leg details for a particular ticket number.

```
SELECT * FROM
NESTED TABLES (ticket.bagInfo.flightLegs b ancestors(ticket a))
WHERE a.ticketNo=1762344493810
```

Explanation: This is an example of a join where the target table`ticket` is joined with its descendant `bagInfo`. Additionally, you can limit the result set by applying a filter condition to the result of the join. You are limiting the result to a particular ticket number.

The result has two rows, implying there are two flight legs for this ticket number.

Output:

```
"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM604",
"flightDate":"2019-02-01T06:00:00.0000Z","fltRouteSrc":"MIA","fltRouteDest":"L
AX",
"estimatedArrival":"2019-02-01T11:00:00.0000Z",
"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"}]}}

{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},
"b":{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM667",
"flightDate":"2019-02-01T06:13:00.0000Z","fltRouteSrc":"LAX","fltRouteDest":"M
EL",
"estimatedArrival":"2019-02-01T16:15:00.0000Z",
"actions":[{"actionAt":"MEL","actionCode":"Offload to Carousel at
MEL","actionTime":"2019-02-01T16:15:00Z"},
{"actionAt":"LAX","actionCode":"ONLOAD to
MEL","actionTime":"2019-02-01T15:35:00Z"},
{"actionAt":"LAX","actionCode":"OFFLOAD from
LAX","actionTime":"2019-02-01T15:18:00Z"}]}}
```

Example 4: Fetch the bag id and number of hops for all bags of all passengers.

```
SELECT b.id,count(*) AS NUMBER_HOPS
FROM NESTED TABLES (ticket a descendants(ticket.bagInfo.flightLegs b))
GROUP BY b.id
```

Explanation: You group the data based on the bag id (using GROUP BY) and get the count of flight legs (using count()) for every bag.

Output:

```
{"id":79039899168383,"NUMBER_HOPS":3}
{"id":79039899156435,"NUMBER_HOPS":2}
{"id":7903989918469,"NUMBER_HOPS":1}
{"id":79039899165297,"NUMBER_HOPS":2}
{"id":79039899197492,"NUMBER_HOPS":3}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 4a: Find the number of hops for all bags of a particular passenger.

```
SELECT b.id,count(*) AS NUMBER_HOPS FROM
NESTED TABLES (ticket a descendants(ticket.bagInfo.flightLegs b))
WHERE a.ticketNo=1762355527825
GROUP BY b.id
```

Explanation: You group the data based on the bag id (using GROUP BY) and get the count of flight legs (Using count()) for every bag. Additionally, you filter the results for a particular ticket number.

Output:

```
{"id":79039899197492,"NUMBER_HOPS":3}
```

Example 5: Fetch bag id and routing details of all bags that arrived after 2019.

```
SELECT b.id, routing FROM
NESTED TABLES(ticket a descendants(ticket.bagInfo b))
WHERE CAST (b.bagArrivalDate AS Timestamp(0))>=
CAST ("2019-01-01T00:00:00" AS Timestamp(0))
```

Explanation: This is an example of a join where the target table`ticket` is joined with its child table `bagInfo`. The filter condition is applied on the `bagArrivalDate`. The `CAST` function is used to convert the string into Timestamp and then the values are compared.

Output:

```
{"id":79039899197492,"routing":"BZN/SEA/CDG/MXP"}
{"id":79039899165297,"routing":"MIA/LAX/MEL"}
{"id":79039899168383,"routing":"MXP/CDG/SLC/BZN"}
{"id":79039899156435,"routing":"GRU/ORD/SEA"}
{"id":7903989918469,"routing":"JFK/MAD"}
```

Query API examples

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***TableJoins.java*** from the examples here.

```
   /* fetch rows based on joins*/
private static void fetchRows(NoSQLHandle handle,String sql_stmt) throws 
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sql_stmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)) {
         System.out.println("Query results:");
         for (MapValue res : results) {
            System.out.println("\t" + res);
         }
      }
}

System.out.println("Fetching data using NESTED TABLES:");
String sql_stmt_nt ="SELECT * FROM NESTED TABLES (ticket a 
descendants(ticket.bagInfo.flightLegs b))";
/* fetching rows using nested tables*/
fetchRows(handle,sql_stmt_nt);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***TableJoins.py*** from the examples here.

```
# Fetch data from the table based on joins
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))

sql_stmt_nt='SELECT * FROM NESTED TABLES (ticket a 
descendants(ticket.bagInfo.flightLegs b))'
print('Fetching data using NESTED TABLES ')
fetch_data(handle,sql_stmt_nt)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***TableJoins.go*** from the examples here.

```
func fetchData(client *nosqldb.Client, err error,
               tableName string, querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{ Statement: querystmt,}

   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,}
   var results []*types.MapValue

   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()

      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
      }

      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
   }
   for i, r := range results {
      fmt.Printf("\t%d: %s\n", i+1,
                  jsonutil.AsJSON(r.Map()))
   }
}

querystmt_nt:= "SELECT * FROM NESTED TABLES (ticket a
descendants(ticket.bagInfo.flightLegs b)"
fmt.Println("Fetching data using NESTED TABLES")
fetchData(client, err,querystmt_nt)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***TableJoins.js*** from the examples here.

```
//fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
const stmt_nt = 'SELECT * FROM NESTED TABLES (ticket a
descendants(ticket.bagInfo.flightLegs b))';
console.log("Fetching data using NESTED TABLES");
await fetchData(handle,stmt_nt);
```

TypeScript: Download the full code ***TableJoins.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}

const stmt_nt = 'SELECT * FROM NESTED TABLES (ticket a
descendants(ticket.bagInfo.flightLegs b))';
console.log("Fetching data using NESTED TABLES");
await fetchData(handle,stmt_nt);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***TableJoins.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Row
      {
         Console.WriteLine();
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            Console.WriteLine(row.ToJsonString());
        }
    }
}

private const string stmt_nt ="SELECT * FROM NESTED TABLES (ticket a
descendants(ticket.bagInfo.flightLegs b))";
Console.WriteLine("Fetching data using NESTED TABLES: ");
await fetchData(client,stmt_nt);
```

Using inner join with parent-child tables

A Join is used to combine rows from two or more tables, based on related columns or fields between them. In a hierarchical table, the 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. All tables in the hierarchy have the same shard key columns.

An inner join is one of the types of join used to combine tables that belong to the same table hierarchy in an Oracle NoSQL Database.

- Overview of Inner Join
- Examples using Inner Join

Overview of 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.

Conceptually, an inner join works as follows:

Consider that you need to perform an inner join of three tables A, B and C. The tables A and B are first joined. That is, if table A has N rows and n columns, and table B has M rows and m columns, every row in table A is joined with every row in table B. The resultant table AB would thus have (N * M) rows and (n + m) columns. Similarly table AB is now joined with table C, to form table ABC. The join predicates in the WHERE clause are then applied to the table ABC. Note that the join predicates must include equality predicates between all the shard key columns of the joined tables. The final result-set contains only the matching rows from the participating tables.

You specify the tables to be joined in the FROM clause of the SELECT statement and the join predicates in the WHERE clause. A join predicate is a predicate that references the columns or fields from one or more tables that are to be joined and specifies the filter conditions that need to be applied on them. In the case of inner join, the WHERE clause must include the equality predicate on all shard keys of the participating tables.

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 in the result-set is as specified in the SELECT clause.

While performing an inner join, the following are applicable:

- Only joins among tables in the same table hierarchy are allowed.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- 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. To know more about shard keys, see CREATE TABLE. That is, for any pair of joined tables, a row from one table matches with a row from the other table only if they both have the same values on their shard key columns. 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.

An inner join differs from NESTED TABLES and left outer join primarily in the following aspects:

- An inner join is based on matching the shard keys of the participating tables, whereas NESTED TABLES and left outer join are based on matching the primary keys of the participating tables.
- The result-set of an inner join contains only the matching rows. Whereas, in the case of NESTED TABLES and left outer join, the unmatched row in the left table is also returned in the result-set with a corresponding NULL row in the right table.
- Inner join can be used to join tables that are not in an ancestor-descendant relationship. This is not possible in the case of left outer join and NESTED_TABLES. For more details, see Inner Join vs LOJ vs NESTED TABLES.

In essence, tables having an 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.

Examples using Inner Join

Consider an airline baggage tracking application. 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 `flightlegs` which is the child of the `baggageInfo` table.

Download the script parentchildtbls_loaddata.sql and run it as shown below. This script creates the tables used in the example and loads data into the tables.

- Start your KVSTORE or KVLite

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

- Open the SQL shell

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

  The SQL prompt appears.
- Load the DDL file to create the necessary tables used in the example

  ```
  load -file parentchild.ddl
  ```


Copyright © 2022, 2025, Oracle and/or its affiliates.

- Use the `load` command to run the script. 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 ###
```

Following are the tables created:

- **`ticket`**

  ```
      ticketNo LONG
      confNo STRING
      PRIMARY KEY(ticketNo)
  ```

- **`ticket.bagInfo`**

  ```
      id LONG
      tagNum LONG
      routing STRING
      lastActionCode STRING
      lastActionDesc STRING
      lastSeenStation STRING
      lastSeenTimeGmt TIMESTAMP(4)
      bagArrivalDate TIMESTAMP(4)
      PRIMARY KEY(id)
  ```

- **`ticket.bagInfo.flightLegs`**

  ```
      flightNo STRING
      flightDate TIMESTAMP(4)
      fltRouteSrc STRING
      fltRouteDest STRING
      estimatedArrival TIMESTAMP(4)
      actions JSON
      PRIMARY KEY(flightNo)
  ```

- **`ticket.passengerInfo`**

  ```
      contactPhone STRING
      fullName STRING
  ```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
gender STRING
PRIMARY KEY(contactPhone)
```

SQL Examples

Let us now see a few example SQL queries for inner join:

Example 1: Fetch the details of the passenger with ticket number 1762324912391.

```
SELECT fullname, contactPhone, gender FROM ticket a,ticket.passengerInfo b
WHERE
      a.ticketNo=b.ticketNo AND a.ticketNo=1762324912391
```

Explanation: This is an example of an inner join where the parent table ticket is joined with its child table passengerInfo and a filter is applied to restrict the result. Note that the shard key here is ticketNo. If the shard key is not explicitly specified while creating the root table, the primary key of the root table is taken as the shard key. This shard key is inherited by all the descendant tables.

Output:

```
{"fullname":"Elane Lemons","contactPhone":"600-918-8404","gender":"F"}
1 row returned
```

Example 2: Fetch the bag details of all passengers who have been issued a ticket.

```
SELECT * FROM ticket a, ticket.bagInfo b WHERE a.ticketNo=b.ticketNo
```

Explanation: This is an example of an inner join where the parent table `ticket` is joined with its child table `bagInfo`.

Output:


```
{"a":{"ticketNo":1762324912391,"confNo":"LN0C8R"},"b":
{"ticketNo":1762324912391,"id":79039899168383,"tagNum":1765780623244,"routing"
:"MXP/CDG/SLC/
BZN","lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"
BZN","lastSeenTimeGmt":"2019-03-15T10:13:00.0000Z","bagArrivalDate":"2019-03-1
5T10:13:00.0000Z"}}
{"a":{"ticketNo":1762355527825,"confNo":"HJ4J4P"},"b":
{"ticketNo":1762355527825,"id":79039899197492,"tagNum":17657806232501,"routing
":"BZN/SEA/CDG/
MXP","lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"
MXP","lastSeenTimeGmt":"2019-03-22T10:17:00.0000Z","bagArrivalDate":"2019-03-2
2T10:17:00.0000Z"}}
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},"b":
{"ticketNo":1762344493810,"id":79039899165297,"tagNum":17657806255240,"routing
":"MIA/LAX/
MEL","lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"
MEL","lastSeenTimeGmt":"2019-02-01T16:13:00.0000Z","bagArrivalDate":"2019-02-0
1T16:13:00.0000Z"}}
{"a":{"ticketNo":1762376407826,"confNo":"ZG8Z5N"},"b":
{"ticketNo":1762376407826,"id":7903989918469,"tagNum":17657806240229,"routing"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
:"JFK/
MAD","lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"
MAD","lastSeenTimeGmt":"2019-03-07T13:51:00.0000Z","bagArrivalDate":"2019-03-0
7T13:51:00.0000Z"}}
{"a":{"ticketNo":1762392135540,"confNo":"DN3I4Q"},"b":
{"ticketNo":1762392135540,"id":79039899156435,"tagNum":17657806224224,"routing
":"GRU/ORD/
SEA","lastActionCode":"OFFLOAD","lastActionDesc":"OFFLOAD","lastSeenStation":"
SEA","lastSeenTimeGmt":"2019-02-15T21:21:00.0000Z","bagArrivalDate":"2019-02-1
5T21:21:00.0000Z"}}
5 rows returned
```


Example 3: Fetch the flight leg details of the bags of the passenger with ticket number 1762344493810.

```
SELECT * FROM ticket a, ticket.bagInfo.flightLegs b WHERE
a.ticketNo=b.ticketNo AND
      a.ticketNo=1762344493810
```

Explanation: : This is an example of an inner join where the parent table `ticket` is joined with its descendant `flightlegs`. A descendant table can be any level hierarchically below a table (For example `flightLegs` is the child of `bagInfo` which is the child of `ticket`, so `flightLegs` is a descendant of `ticket`). The result is then filtered for a particular ticket number.

Output:


```
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},"b":
{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM604","flightDate":
"2019-02-01T06:00:00.0000Z","fltRouteSrc":"MIA","fltRouteDest":"LAX","estimate
dArrival":"2019-02-01T11:00:00.0000Z","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"}]}}
{"a":{"ticketNo":1762344493810,"confNo":"LE6J4Z"},"b":
{"ticketNo":1762344493810,"id":79039899165297,"flightNo":"BM667","flightDate":
"2019-02-01T06:13:00.0000Z","fltRouteSrc":"LAX","fltRouteDest":"MEL","estimate
dArrival":"2019-02-01T16:15:00.0000Z","actions":
[{"actionAt":"MEL","actionCode":"Offload to Carousel at
MEL","actionTime":"2019-02-01T16:15:00Z"},
{"actionAt":"LAX","actionCode":"ONLOAD to
MEL","actionTime":"2019-02-01T15:35:00Z"},
{"actionAt":"LAX","actionCode":"OFFLOAD from
LAX","actionTime":"2019-02-01T15:18:00Z"}]}}
2 rows returned
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 4: Find the number of hops for all the bags of a passenger with ticket number 1762355527825 . If there are multiple bags checked in for a passenger, then the number of hops for all the bags are displayed.

```
SELECT b.id,count(*) AS NUMBER_HOPS FROM ticket a, ticket.bagInfo.flightLegs
b WHERE a.ticketNo=b.ticketNo AND a.ticketNo=1762355527825 GROUP BY
    b.id
```

Explanation: Here, you group the data based on the bag id (using GROUP BY) and get the count of flight legs (using count()) for every bag. Additionally, you filter the results for a particular ticket number.

Output:

```
{"id":79039899197492,"NUMBER_HOPS":3}
1 row returned
```

Example 5: 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: This is an example of an inner join of three tables, that is, the parent table `ticket`, and the sibling tables `passengerInfo` and `bagInfo`.

Output:

```
{"ticketNo":1762324912391,"fullName":"Elane
Lemons","bagArrivalDate":"2019-03-15T10:13:00.0000Z"}
{"ticketNo":1762355527825,"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: Fetch the name of the passenger, the last seen station of whose bag is “MEL”

```
SELECT a.fullName FROM ticket.passengerInfo a, ticket.bagInfo b WHERE
a.ticketNo = b.ticketNo AND b.lastSeenStation = "MEL"
```

Explanation: This is an example of an inner join of the sibling tables `passengerInfo` and `bagInfo`. The name of the passenger whose bag was last seen at the "MEL" station is returned.

Output:

```
{"fullName":"Adam Phillips"}
1 row returned
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 7: 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: This is an inner join of two tables, `passengerInfo` and `flightlegs`, that are not in an ancestor-descendant relationship.

Output:

```
{"fullName":"Adam Phillips"}
```

Such a join between tables that are not in an ancestor-descendant relationship is not possible with Left Outer Join and NESTED TABLES.

Query API Examples

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***TableJoins.java*** from the examples here.

```
/* fetch rows based on joins*/
private static void fetchRows(NoSQLHandle handle,String sql_stmt) throws
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sql_stmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)) {
         System.out.println("Query results:");
         for (MapValue res : results) {
            System.out.println("\t" + res);
         }
      }
}

/* fetching rows using inner join*/
String sql_stmt_innerjoin ="SELECT * FROM ticket a, ticket.bagInfo.flightLegs
b WHERE a.ticketNo=b.ticketNo";
System.out.println("Fetching data using inner join:");
fetchRows(handle,sql_stmt_innerjoin);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***TableJoins.py*** from the examples here

```
# Fetch data from the table based on joins
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))

sql_stmt_ij='SELECT * FROM ticket a, ticket.bagInfo.flightLegs b WHERE
a.ticketNo=b.ticketNo'
print('Fetching data using Inner Join ')
fetch_data(handle,sql_stmt_ij)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***TableJoins.go*** from the examples here.

```
func fetchData(client *nosqldb.Client, err error,
               tableName string, querystmt string)(){
   prepReq := &nosqldb.PrepareRequest{ Statement: querystmt,}

   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }

   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,}
   var results []*types.MapValue

   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()

      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
      }

      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
    for i, r := range results {
       fmt.Printf("\t%d: %s\n", i+1,
                   jsonutil.AsJSON(r.Map()))
    }
}

querystmt_ij:= "SELECT * FROM ticket a, ticket.bagInfo.flightLegs b WHERE
a.ticketNo=b.ticketNo"
fmt.Println("Fetching data using Inner Join")
fetchData(client, err,querystmt_ij)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***TableJoins.js*** from the examples here.

```
//fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}

const stmt_ij = 'SELECT * FROM ticket a, ticket.bagInfo.flightLegs b WHERE
a.ticketNo=b.ticketNo';
console.log("Fetching data using Inner Join");
await fetchData(handle,stmt_ij);
```

TypeScript: Download the full code ***TableJoins.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            for(let row of result.rows) {
               console.log('  %O', row);
            }
            opt.continuationKey = result.continuationKey;
         } while(opt.continuationKey);
      } catch(error) {
         console.error('  Error: ' + error.message);
      }
   }

const stmt_ij = 'SELECT * FROM ticket a, ticket.bagInfo.flightLegs b WHERE
a.ticketNo=b.ticketNo';
console.log("Fetching data using Inner Join");
await fetchData(handle,stmt_ij);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***TableJoins.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Row)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}

private const string stmt_ij ="SELECT * FROM ticket a,
ticket.bagInfo.flightLegs b WHERE a.ticketNo=b.ticketNo";
Console.WriteLine("Fetching data using Inner Join: ");
await fetchData(client,stmt_ij);
```

Tuning and Optimizing SQL queries

Query optimization is the overall process of choosing the most efficient means of executing a SQL statement.

You optimize a SQL query to get accurate and fast database results.

- Using Indexes for query optimization
- Examples of queries using index

Copyright © 2022, 2025, Oracle and/or its affiliates.

Using Indexes for query optimization

Indexing is a way to optimize the performance of a database by minimizing the number of disk accesses required when a query is processed.

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?

There are no statistics on the number and distribution of values in a table column. 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. You can use index hints to force the use of a particular index in queries.

Examples of queries using index

You can write simple queries to understand how an index is used.

Query 1:

Fetch the bag details of passengers for ticket numbers satisfying 2 range of values.

```
SELECT fullname, ticketNo,bag.bagInfo[].tagNum,
bag.bagInfo[].routing
FROM BaggageInfo bag WHERE 1762340000000 < ticketNo
AND ticketNo < 1762352000000
```

In the above example, the query contains 2 index predicates. The primary key index is used as `ticketNo` is the primary key here. For the primary key index, 1762340000000 < ticketNo is a start predicate and ticketNo < 1762352000000 is a stop predicate.

A portion of the query plan is shown below. You can see the primary index being used.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "primary index",
"covering index" : false,
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
"index scans" :
[
  {
    "equality conditions" : {},
    "range conditions" : { "ticketNo" : { "start value" : 1762340000000,
       "start inclusive" : false,
       "end value" : 1762352000000,
       "end inclusive" : false } }
  }
]
```

For more information on how a query is executed, see Query execution plan.

Query 2:

Fetch the bag details of passengers for ticket numbers satisfying one of the two ranges of values.

```
SELECT fullname, ticketNo,bag.bagInfo[].tagNum,
bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE ticketNo > 1762340000000 OR
ticketNo < 1762352000000
```

In the above example, the query contains 1 index predicate, which is the whole WHERE expression. The primary key index is used as `ticketNo` is the primary key here. The predicate is a filtering predicate.

A portion of the query plan is shown below. You can see the primary index and the index filtering predicates being used.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "primary index",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {},
    "range conditions" : {}
  }
],
"index filtering predicate" :
{
  "iterator kind" : "OR",
  "input iterators" :
  [
    {
      "iterator kind" : "GREATER_THAN",
      "left operand" :
        {
          ---
        },
      "right operand" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        {
          ---
        }
    },
    {
      "iterator kind" : "LESS_THAN",
      "left operand" :
        {
          ---
        },
      "right operand" :
        {
          ---
        }
    }
  ]
}
```

For more information on how a query is executed, see Query execution plan.

Query 3:

Fetch the bag details for a particular reservation code.

```
SELECT fullName,bag.ticketNo, bag.confNo,
bag.bagInfo[].tagNum,bag.bagInfo[].routing
FROM BaggageInfo bag WHERE bag.confNo="FH7G1W"
```

In the above example, two indexes are applicable `compindex_tckNoconfNo` and `fixedschema_conf` .

A portion of the query plan is shown below. The `fixedschema_conf` is used as that is a single index on `ticketNo`. An index scan is performed with the equality condition.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "fixedschema_conf",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {"confNo":"FH7G1W"},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Query 4:

Fetch the name and routing details of all male passengers.

```
SELECT fullname,bag.bagInfo[].routing FROM BaggageInfo bag
WHERE gender!="F"
```

In the above example, there is no index predicate, because no index has information about gender.

A portion of the query plan is shown below. As there are no available indexes to be used, only the primary key index is used.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "primary index",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Query 5:

Fetch the name and phone number for all passengers.

```
SELECT bag.contactPhone, bag.fullName FROM BaggageInfo bag
ORDER BY bag.fullName
```

In the above example, only the index `compindex_namephone` is applicable. The sort (for the order by clause) will be index-based because the order-by expression matches the 1st field of the index used by the query. In this case, the full name and contact phone 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. So the index `compindex_namephone` is a covering index in this example. The query processor will apply this optimization.

A portion of the query plan is shown below. You can see the index `compindex_namephone` is used and it is a covering index.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "compindex_namephone",
"covering index" : true,
"index row variable" : "$$bag_idx",
"index scans" :
[
  {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    "equality conditions" : {},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Query 6:

Fetch the name, ticket number, and arrival date of passengers whose arrival date is greater than a given value.

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

In the above example, the EXISTS condition is actually converted to a filtering predicate. There is one filtering predicate which is the whole WHERE expression.

A portion of the query plan is shown below. The index `simpleindex_arrival` is used in this example.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "simpleindex_arrival",
"covering index" : false,
"index row variable" : "$$bag_idx",
"index scans" :
[
  {
    "equality conditions" : {},
    "range conditions" : {}
  }
],
"index filtering predicate" :
{
  "iterator kind" : "GREATER_OR_EQUAL",
  "left operand" :
    {
      ---
    },
  "right operand" :
    {
      ---
    }
}
```

For more information on how a query is executed, see Query execution plan.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Query 7:

Fetch the reservation code and count of bags for all passengers.

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

In the above example, two indexes `fixedschema_conf` and `compindex_tckNoconfNo` are applicable.

A portion of the query plan is shown below. The index `fixedschema_conf` is used as that is a single index with only one column `confNo`. For this query, the group-by is index-based. As you need the entire bagInfo details to determine the number of bags using the aggregate **count** function, the index here is not covering.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "fixedschema_conf",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Query 8:

Fetch the full name and tag number of passengers who are in the given list of names.

```
SELECT bagdet.fullName, bagdet.bagInfo[].tagNum
FROM BaggageInfo bagdet
WHERE bagdet.fullName IN
("Lucinda Beckman", "Adam Phillips",
"Zina Christenson","Fallon Clements")
```

In the above example, only the index `compindex_namephone` is applicable.

A portion of the query plan is shown below. The index `compindex_namephone` is used. An index scan is performed on `compindex_namephone` evaluating four equality predicates.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bagdet",
"index used" : "compindex_namephone",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {"fullName":"Lucinda Beckman"},
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    "range conditions" : {}
  },
  {
    "equality conditions" : {"fullName":"Adam Phillips"},
    "range conditions" : {}
  },
  {
    "equality conditions" : {"fullName":"Zina Christenson"},
    "range conditions" : {}
  },
  {
    "equality conditions" : {"fullName":"Fallon Clements"},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Query 9:

Select the ticket details(ticket number, reservation code, tag number, and routing) for a passenger with a specific ticket number and reservation code.

```
SELECT fullName,bag.ticketNo, bag.confNo,
bag.bagInfo[].tagNum,bag.bagInfo[].routing
FROM BaggageInfo bag WHERE
bag.ticketNo=1762311547917
AND bag.confNo="FH7G1W"
```

In the above example, though the index `compindex_tckNoconfNo` is available, only the primary index (for `ticketNo`) gets used. An index scan is performed on the primary index and the WHERE expression is evaluated.

A portion of the query plan is shown below.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$$bag",
"index used" : "primary index",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {"ticketNo":1762311547917},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Query 10:

Fetch the source of passenger bags and the count of bags for all passengers and group the data by the source.

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

In the above example, there is no index on the `fltRouteSrc` field. So the grouping is done in a generic way. An internal variable is created that iterates over the records produced by the SELECT statement.

A portion of the query plan is shown below. The primary index is being used.

```
"iterator kind" : "TABLE",
"target table" : "BaggageInfo",
"row variable" : "$bag",
"index used" : "primary index",
"covering index" : false,
"index scans" :
[
  {
    "equality conditions" : {},
    "range conditions" : {}
  }
]
```

For more information on how a query is executed, see Query execution plan.

Managing GeoJSON data

The GeoJson specification defines the structure and content of json objects that are supposed to represent geographical shapes on earth (called geometries).

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. See About GeoJSON Data for more details on the various types of geometry objects. 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. See Lines and Coordinate System for more details.

If you want to follow along with the examples, download the script geojsonschema_loaddata.sql and run it as shown below. This script creates the table used in the example and loads data into the table.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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
```

The `geojsonschema_loaddata.sql` contains the following:

```
### Begin Script ###
load -file geoschema.ddl
import -table PointsOfInterest  -file geoschema.json
### End Script ###
```

Using the `load` command, run the script.

```
load -file geojsonschema_loaddata.sql
```

Oracle NoSQL Database implements a number of functions that interpret JSON objects as geometries and allow for the search for rows containing geometries that satisfy certain conditions.

- geo_inside
- geo_intersect
- geo_distance
- geo_within_distance
- geo_near
- geo_is_geometry

geo_inside

Determines geometries within a bounding GeoJSON geometry.

```
boolean geo_inside(any*, any*)
```

- The first parameter `any*` can be any geometric object.
- The second parameter `any*` needs to be a polygon.

The function determines if the geometry pointed by the first parameter is completely contained inside the polygon pointed by the second parameter.

If any of the two parameters does not return a single valid geometry object, and if it can be detected at compile time then the function raises an error.

The runtime behavior is as follows:

- Returns false if any parameter returns 0 or more than 1 item.
- Returns NULL if any parameter returns NULL.
- Returns false if any parameter (at runtime) returns an item that is not a valid geometry object.
- Returns false if the second parameter returns a geometry object that is not a polygon.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- If both parameters return a single geometry object each 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.
  - Else it returns false.

> **Note**
>
> The interior of a polygon is all the points in the polygon area except the points on the linear ring that define the polygon's boundary.

Example: Look for nature parks in Northern California.

```
SELECT t.poi.name AS park_name,
t.poi.address.street AS park_location
FROM PointsOfInterest t
WHERE t.poi.kind = "nature park"
AND geo_inside(t.poi.location,
              { "type" : "polygon",
                "coordinates": [[
                  [-120.1135253906249, 36.99816565700228],
                  [-119.0972900390625, 37.391981943533544],
                  [-119.2840576171875, 37.97451499202459],
                  [-120.2069091796874, 38.035112420612975],
                  [-122.3822021484375, 37.74031329210266],
                  [-122.2283935546875, 37.15156050223665],
                  [-121.5362548828124, 36.85325222344018],
                  [-120.1135253906249, 36.99816565700228]
                ]]
             });
```

Explanation:

- You query the `PointsOfInterest` table to filter the rows for **nature park**.
- You specify a polygon as the second parameter to the `geo_inside` function.
- The coordinates of the polygon you specify correspond to the coordinates of the northern portion of the state of California in the U.S.
- The `geo_inside` function only returns rows when the location of the nature park is completely contained inside the location points specified.

Result:

```
{"park_name":"portola redwoods state park",
"park_location":"15000 Skyline Blvd"}
```

geo_intersect

Determines geometries that intersect with a GeoJSON geometry.

```
boolean geo_intersect(any*, any*)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

The first and the second parameters `any*` can be any geometric object.

The function determines if two geometries that are specified as parameters have any points in common. If any of the two parameters does not return a single valid geometry object, and if it can be detected at compile time then the function raises an error.

The runtime behavior is as follows:

- Returns false if any parameter returns 0 or more than 1 item.
- Returns NULL if any parameter returns NULL.
- Returns false if any parameter (at runtime) returns an item that is not a valid geometry object.

If both parameters return a single geometry object each, the function returns true if the 2 geometries have any points in common; otherwise false.

Example: Texas is considering regulating access to the underground water supply. An aquifer is an underground layer of water-bearing permeable rock, rock fractures, or unconsolidated materials. The government wants to impose new regulations for locations that are very close to an aquifer.

The coordinates of the aquifer have already been mapped. You want to know all counties in the Texas state that intersect with that aquifer so that you can notify the county government for each affected county to participate in talks for the new regulations.

```
SELECT t.poi.county AS County_needs_regulation,
t.poi.contact AS Contact_phone
FROM PointsOfInterest t WHERE
geo_intersect(
    t.poi.location,
    {

     "type" : "polygon",
      "coordinates": [
          [
            [-97.668457031249, 29.34387539941801],
            [-95.207519531258, 29.19053283229458],
            [-92.900390625653, 30.37287518811801],
            [-94.636230468752, 32.21280106801518],
            [-97.778320312522, 32.45415593941475],
            [-99.799804687541, 31.18460913574325],
            [-97.668457031249, 29.34387539941801]
          ]
        ]
    }
);
```

Explanation:

- The above query fetches the locations which intersect with the location of the aquifer. That is if the location coordinates have any points in common with the location of the aquifer.
- You use `geo_intersect` to see if the coordinates of the location have any points common with the coordinates of the aquifer that are specified.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Result:

```
{"County_needs_regulation":"Tarrant","Contact_phone":"469 745 5687"}
{"County_needs_regulation":"Kinga","Contact_phone":"469 384 7612"}
```

geo_distance

Determines distance between two geospatial objects.

```
double geo_distance(any*, any*)
```

The first and the second parameters `any*` can be any geometric object.

The function returns the geodetic distance between the two 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.

Overview of Geodetic Line

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's surface is a sphere. Then, the geodetic line between two points on the 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.

If any of the two parameters does not return a single valid geometry object, and if it can be detected at compile time then the function raises an error.

The runtime behavior is as follows:

- Returns -1 if any parameter returns zero or more than 1 item.
- Returns NULL if any parameter returns NULL.
- Returns -1 if any of the parameters is not a geometry object.

Otherwise, the function returns the geodetic distance in meters between the 2 input geometries.

Copyright © 2022, 2025, Oracle and/or its affiliates.

> **Note**
>
> The results are sorted ascending by distance(displaying the shortest distance first).

Example: How far is the nearest restaurant from the given location?

```
SELECT
t.poi.name AS restaurant_name,
t.poi.address.street AS street_name,
geo_distance(
    t.poi.location,
    {
       "type" : "point",
       "coordinates": [-121.94034576416016,37.2812239247177]
    }
) AS distance_in_meters
FROM PointsOfInterest t
WHERE t.poi.kind = "restaurant" ;
```

Explanation:

- You query the `PointsOfInterest` table to filter the rows for **restaurant**.
- You provide the correct location point and determine the distance using the `geo_distance` function.

Result:

```
{"restaurant_name":"Coach Sports Bar & Grill","street_name":"80 Edward
St","distance_in_meters":799.2645323337218}
{"restaurant_name":"Ricos Taco","street_name":"80 East Boulevard
St","distance_in_meters":976.5361117138553}
{"restaurant_name":"Effie's Restaurant and Bar","street_name":"80 Woodeard
St","distance_in_meters":2891.0508307646282}
```

The distance between the current location and the nearest restaurant is 799 meters.

geo_within_distance

Determines geospatial objects in proximity to a point.

```
boolean geo_within_distance(any*, any*,double)
```

The first and the second parameters `any*` can be any geometric object.

The function determines if the first geometry is within a distance of N meters from the second geometry.

If any of the two parameters does not return a single valid geometry object, and if it can be detected at compile time then the function raises an error.

The runtime behavior is as follows:

- Returns false if any parameter returns 0 or more than 1 item.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- Returns NULL if any of the first two parameters returns NULL.
- Returns false if any of the first two parameters returns an item that is not a valid geometry object.

Finally, if both the parameters return a single geometry object each, 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 parameter; 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.

Example: Is a city hall there in the next 5 km? How far is it?

```
SELECT t.poi.address.street AS city_hall_address,
geo_distance(
    t.poi.location,
    {
        "type" : "point",
        "coordinates" : [-120.653828125,38.85682013474361]
    }

) AS distance_in_meters
FROM PointsOfInterest t
WHERE t.poi.kind = "city hall" AND
geo_within_distance(
    t.poi.location,
    {
        "type" : "point",
        "coordinates" : [-120.653828125,38.85682013474361]
    },
    5000
);
```

Explanation:

- You query the `PointsOfInterest` table to filter the rows for **city hall**.
- You use the `geo_within_distance` function to filter city hall within 5 km (5000m) of the given location.
- You also fetch the actual distance between your location and the city hall using the `geo_distance` function.

Result:

```
{"city_hall_address":"70 North 1st
street","distance_in_meters":1736.0144040331768}
```

The city hall is 1736 m(1.73 km) from the current location.

geo_near

Determines geospatial objects in proximity to a point.

```
boolean geo_near(any*, any*, double)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

The first and the second parameters `any*` can be any geometric object.

The function determines if the first geometry is within a distance of N meters from the second geometry.

If any of the two parameters does not return a single valid geometry object, and if it can be detected at compile time then the function raises an error.

The runtime behavior is as follows:

- Returns false if any parameter returns 0 or more than 1 item.
- Returns NULL if any of the first two parameters returns NULL.
- Returns false if any of the first two parameters returns an item that is not a valid geometry object.

Finally, if both of the first two parameters return a single geometry object each, 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 parameter; otherwise false.

> **Note**
>
> 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.

Example 1: Is there a hospital within 3km of the given location?

```
SELECT
t.poi.name AS hospital_name,
t.poi.address.street AS hospital_address
FROM PointsOfInterest t
WHERE t.poi.kind = "hospital"
AND
geo_near(
    t.poi.location,
    {"type" : "point",
     "coordinates" : [-122.03493933105469,37.32949164059004]
    },
    3000
);
```

Explanation:

- You query the `PointsOfInterest` table to filter the rows for **hospital**.
- You use the `geo_near` function to filter hospitals within 3000m of the given location.

Result:

```
{"hospital_name":"St. Marthas hospital","hospital_address":"18000 West Blvd"}
{"hospital_name":"Memorial hospital","hospital_address":"10500 South St"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 2: How far is a gas station within the next one mile from the given location?

```
SELECT
t.poi.address.street AS gas_station_address,
geo_distance(
    t.poi.location,
    {
        "type" : "point",
        "coordinates" : [-121.90768646240233,37.292081740702365]
    }
) AS distance_in_meters
FROM PointsOfInterest t
WHERE t.poi.kind = "gas station" AND
geo_near(
    t.poi.location,
    {
        "type" : "point",
        "coordinates" : [-121.90768646240233,37.292081740702365]
    },
    1600
);
```

Explanation:

- You query the `PointsOfInterest` table to filter the rows for **gas station**.
- You use the `geo_near` function to filter gas stations within one mile(1600m) of the given location.
- You also fetch the actual distance between your location and the gas station using the `geo_distance` function.

Result:

```
{"gas_station_address":"33 North
Avenue","distance_in_meters":886.7004173859665}
```

The actual distance to the nearest gas station within the next mile is 886m.

geo_is_geometry

Validates a geospatial object.

```
boolean geo_is_geometry(any*)
```

The parameter `any*` can be any geometric object.

The function determines if the given input is a valid geometry object.

- Returns false if the parameter returns zero or more than 1 item.
- Returns NULL if the parameter returns NULL.
- Returns true if the input is a single valid geometry object. Otherwise, false.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Example: Determine if the location pointing to the **city hall** is a valid geometric object.

```
SELECT geo_is_geometry(t.poi.location) AS city_hall
FROM PointsOfInterest t
WHERE t.poi.kind = "city hall"
```

Explanation: You use the function geo_is_geometry to determine if a given location is a valid geometric object or not.

Result:

```
{ "city_hall" : true}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

5 Reference

The articles in this section contain reference information related to various operators, constructs and expressions used in SQL.

Operators in SQL

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Sequence Comparison Operators
- Logical operators
- NULL operators
- Value Comparison Operators
- IN Operator
- Regular Expression Conditions
- EXISTS Operator
- Is-Of-Type Operator
- SQL Operators examples using QueryRequest API
- BETWEEN Operator

Sequence Comparison Operators

Comparisons between two sequences are done via a set of operators: =any, !=any, >any, >=any, <any, <=any. The result of 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 1: Find passenger name and tag number for all bags where the estimated arrival time is greater than **2019-03-01T13:00:00Z**.

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

Explanation: You fetch the full name, and tag number of all passenger bags whose estimated arrival time is greater than the given value. Here the operand on the left hand of the ">" operator (`bag.bagInfo[].flightLegs[].estimatedArrival`) is a sequence of values. If you try

Copyright © 2022, 2025, Oracle and/or its affiliates.

using the regular comparison operator instead of the sequence operator, you get an error as shown below. That is the reason you need a sequence operator here.

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

Output showing error:

```
Error handling command SELECT fullname,
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].estimatedArrival
FROM BaggageInfo bag WHERE bag.bagInfo[].flightLegs[].estimatedArrival >
"2019-03-01T13:00:00Z":
Error: at (1, 107) 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.
```

Output (after using sequence operator):

```
{"fullname":"Lucinda Beckman","tagNum":"17657806240001","estimatedArrival":
["2019-03-12T16:00:00Z","2019-03-13T03:14:00Z","2019-03-12T15:12:00Z"]}
{"fullname":"Elane Lemons","tagNum":"1765780623244","estimatedArrival":
["2019-03-15T09:00:00Z","2019-03-15T10:14:00Z","2019-03-15T10:14:00Z"]}
{"fullname":"Dierdre
Amador","tagNum":"17657806240229","estimatedArrival":"2019-03-07T14:00:00Z"}
{"fullname":"Henry Jenkins","tagNum":"17657806216554","estimatedArrival":
["2019-03-02T09:00:00Z","2019-03-02T13:24:00Z"]}
{"fullname":"Lorenzo Phil","tagNum":
["17657806240001","17657806340001"],"estimatedArrival":
["2019-03-12T16:00:00Z","2019-03-13T03:14:00Z",
"2019-03-12T15:12:00Z","2019-03-12T16:40:00Z","2019-03-13T03:18:00Z","2019-03-
12T15:12:00Z"]}
{"fullname":"Gerard Greene","tagNum":"1765780626568","estimatedArrival":
["2019-03-07T17:00:00Z","2019-03-08T04:10:00Z","2019-03-07T16:10:00Z"]}
{"fullname":"Doris Martin","tagNum":"17657806232501","estimatedArrival":
["2019-03-22T09:00:00Z","2019-03-21T23:24:00Z","2019-03-22T10:24:00Z"]}
{"fullname":"Omar Harvey","tagNum":"17657806234185","estimatedArrival":
["2019-03-02T02:00:00Z","2019-03-02T16:21:00Z"]}
{"fullname":"Mary Watson","tagNum":"17657806299833","estimatedArrival":
["2019-03-13T15:00:00Z","2019-03-14T06:22:00Z"]}
{"fullname":"Kendal Biddle","tagNum":"17657806296887","estimatedArrival":
["2019-03-04T22:00:00Z","2019-03-05T12:02:00Z"]}
```

Example 2: Find the tag number of passengers who fly from JFK/through JFK to any other location.

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

Explanation: You fetch the tag number of passengers whose flight source is JFK or the passengers who travel through JFK. The destination can be anything.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"tagNum":"17657806240229","fltRouteSrc":"JFK"}
{"tagNum":"17657806215913","fltRouteSrc":["JFK","IST"]}
{"tagNum":"17657806296887","fltRouteSrc":["JFK","IST"]}
```

Logical operators

The operators AND and OR are binary and the NOT operator is unary. The operands of the logical operators are conditional expressions, which must have a type `BOOLEAN`. An empty result from an operand is treated as a false value. If an operand returns NULL(either SQL NULL or JSON 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 1: Select the details of the passenger and their bags for a trip with ticket number **1762311547917** or confirmation number **KN4D1L**.

```
SELECT fullName,bag.ticketNo, bag.confNo,
bag.bagInfo[].tagNum, bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE bag.ticketNo=1762311547917 OR bag.confNo="KN4D1L"
```

Explanation: You fetch the details of passengers satisfying one of the two filter criteria. You do this with the **OR** clause. You fetch the full name, tag number, ticket number, reservation code, and routing details of passengers satisfying a particular ticket number or a particular reservation code (`confNo`).

Output:

```
{"fullName":"Rosalia
Triplett","ticketNo":1762311547917,"confNo":"FH7G1W","tagNum":"17657806215913"
,"routing":"JFK/IST/VIE"}
{"fullName":"Mary
Watson","ticketNo":1762340683564,"confNo":"KN4D1L","tagNum":"17657806299833","
routing":"YYZ/HKG/BLR"}
```

Example 2: Select baggage details of passengers traveling between **MIA** and **MEL**.

```
SELECT fullName, bag.bagInfo[].tagNum, bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE bag.bagInfo[].flightLegs[].fltRouteSrc =any "MIA" AND
bag.bagInfo[].flightLegs[].fltRouteDest=any "MEL"
```

Explanation: You fetch the details of the passengers traveling between MIA and MEL. Since you need to match 2 conditions here, the flight source and the flight destination, you are using an **AND** operator. Here the flight source could be the starting point of the flight or any transit airport. Similarly, the flight destination could be a transit airport or a final destination.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"fullName":"Zulema Martindale","tagNum":"17657806288937","routing":"MIA/LAX/
MEL"}
{"fullName":"Adam Phillips","tagNum":"17657806255240","routing":"MIA/LAX/MEL"}
{"fullName":"Joanne Diaz","tagNum":"17657806292518","routing":"MIA/LAX/MEL"}
{"fullName":"Zina Christenson","tagNum":"17657806228676","routing":"MIA/LAX/
MEL"}
```

Example 3: Select details of those bags which does not originate from MIA/pass through MIA.

```
SELECT fullName, bag.bagInfo[].tagNum, bag.bagInfo[].routing,
bag.bagInfo[].flightLegs[].fltRouteSrc
FROM BaggageInfo bag
WHERE NOT bag.bagInfo[].flightLegs[].fltRouteSrc=any "MIA"
```

Explanation: You fetch the details of passengers not originating from a particular source. To fetch these details, you are using the **NOT** operator here. You want to fetch details of bags which did not start/go through **MIA**.

Output:

```
{"fullName":"Kendal Biddle","tagNum":"17657806296887","routing":"JFK/IST/
VIE","fltRouteSrc":"JFK"}
{"fullName":"Lucinda
Beckman","tagNum":"17657806240001","routing":"SFO/IST/ATH/
JTR","fltRouteSrc":"SFO"}
{"fullName":"Adelaide Willard","tagNum":"17657806224224","routing":"GRU/ORD/
SEA","fltRouteSrc":"GRU"}
{"fullName":"Raymond Griffin","tagNum":"17657806243578","routing":"MSQ/FRA/
HKG","fltRouteSrc":"MSQ"}
{"fullName":"Elane Lemons","tagNum":"1765780623244","routing":"MXP/CDG/SLC/
BZN","fltRouteSrc":"MXP"}
{"fullName":"Dierdre Amador","tagNum":"17657806240229","routing":"JFK/
MAD","fltRouteSrc":"JFK"}
{"fullName":"Henry Jenkins","tagNum":"17657806216554","routing":"SFO/ORD/
FRA","fltRouteSrc":"SFO"}
{"fullName":"Rosalia Triplett","tagNum":"17657806215913","routing":"JFK/IST/
VIE","fltRouteSrc":"JFK"}
{"fullName":"Lorenzo Phil","tagNum":
["17657806240001","17657806340001"],"routing":["SFO/IST/ATH/JTR","SFO/IST/ATH/
JTR"],"fltRouteSrc":["SFO","SFO"]}
{"fullName":"Gerard Greene","tagNum":"1765780626568","routing":"SFO/IST/ATH/
JTR","fltRouteSrc":"SFO"}
{"fullName":"Doris Martin","tagNum":"17657806232501","routing":"BZN/SEA/CDG/
MXP","fltRouteSrc":"BZN"}
{"fullName":"Omar Harvey","tagNum":"17657806234185","routing":"MEL/LAX/
MIA","fltRouteSrc":"MEL"}
{"fullName":"Fallon
Clements","tagNum":"17657806255507","routing":"MXP/CDG/SLC/
BZN","fltRouteSrc":"MXP"}
{"fullName":"Lisbeth Wampler","tagNum":"17657806292229","routing":"LAX/TPE/
SGN","fltRouteSrc":"LAX"}
{"fullName":"Teena Colley","tagNum":"17657806255823","routing":"MSQ/FRA/
HKG","fltRouteSrc":"MSQ"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"fullName":"Michelle Payne","tagNum":"17657806247861","routing":"SFO/IST/ATH/
JTR","fltRouteSrc":"SFO"}
{"fullName":"Mary Watson","tagNum":"17657806299833","routing":"YYZ/HKG/
BLR","fltRouteSrc":"YYZ"}
```

NULL operators

The IS NULL operator tests whether the result of its input expression(either SQL expression or JSON object) 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).

Example 1: Fetch ticket number of passengers whose baggage details are available and is NOT NULL.

```
SELECT ticketNo,fullname FROM BaggageInfo bagdet
WHERE bagdet.bagInfo is NOT NULL
```

Explanation: You fetch the details of passengers who have baggage, which means `bagInfo` JSON is not null.

Output:

```
{"ticketNo":1762357254392,"fullname":"Teena Colley"}
{"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"}
```

Example 2: Fetch ticket number of passengers whose baggage details are not available or IS NULL

```
SELECT ticketNo, fullname FROM BaggageInfo bagdet
WHERE bagdet.bagInfo is NULL
0 row returned
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Value Comparison Operators

Value comparison operators are primarily used to compare 2 values, one produced by the left operand and another from the right operand. If any operand returns more than one item, an error is raised. If both operands return the empty sequence, the operands are considered equal (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 !=. If an operand returns NULL, the result of the comparison expression is also NULL. Otherwise, the result is a boolean value.

Example 1: Select the full name and routing of all male passengers.

```
SELECT fullname, bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE gender="M"
```

Explanation: Here the data is filtered based on gender. The value comparison operator "=" is used to filter the data.

Output:

```
{"fullname":"Lucinda Beckman","routing":"SFO/IST/ATH/JTR"}
{"fullname":"Adelaide Willard","routing":"GRU/ORD/SEA"}
{"fullname":"Raymond Griffin","routing":"MSQ/FRA/HKG"}
{"fullname":"Zina Christenson","routing":"MIA/LAX/MEL"}
{"fullname":"Dierdre Amador","routing":"JFK/MAD"}
{"fullname":"Birgit Naquin","routing":"JFK/MAD"}
{"fullname":"Lorenzo Phil","routing":["SFO/IST/ATH/JTR","SFO/IST/ATH/JTR"]}
{"fullname":"Gerard Greene","routing":"SFO/IST/ATH/JTR"}
{"fullname":"Adam Phillips","routing":"MIA/LAX/MEL"}
{"fullname":"Fallon Clements","routing":"MXP/CDG/SLC/BZN"}
{"fullname":"Lisbeth Wampler","routing":"LAX/TPE/SGN"}
{"fullname":"Teena Colley","routing":"MSQ/FRA/HKG"}
```

You can rewrite this query with a "!=" comparison operator. To get the details of all male passengers, your query can filter data where gender is not "F". This is valid only with the assumption that there can only be two values in the column gender which is "F" and "M".

```
SELECT fullname,bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE gender!="F";
```

Example 2: Fetch the passenger name and routing details of passengers with ticket numbers greater than 1762360000000.

```
SELECT fullname, ticketNo,
bag.bagInfo[].tagNum,bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE ticketNo > 1762360000000
```

Explanation: You need the details of passengers whose ticket number is greater than the given value. You use the ">" operator to filter the data.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"fullname":"Adelaide
Willard","ticketNo":1762392135540,"tagNum":"17657806224224","routing":"GRU/ORD
/SEA"}
{"fullname":"Raymond
Griffin","ticketNo":1762399766476,"tagNum":17657806243578,"routing":"MSQ/FRA/
HKG"}
{"fullname":"Zina
Christenson","ticketNo":1762390789239,"tagNum":"17657806228676","routing":"MIA
/LAX/MEL"}
{"fullname":"Bonnie
Williams","ticketNo":1762397286805,"tagNum":"17657806216554","routing":"SFO/OR
D/FRA"}
{"fullname":"Joanne
Diaz","ticketNo":1762383911861,"tagNum":"17657806292518","routing":"MIA/LAX/
MEL"}
{"fullname":"Kendal
Biddle","ticketNo":1762377974281,"tagNum":"17657806296887","routing":"JFK/IST/
VIE"}
{"fullname":"Dierdre
Amador","ticketNo":1762376407826,"tagNum":"17657806240229","routing":"JFK/
MAD"}
{"fullname":"Birgit
Naquin","ticketNo":1762392196147,"tagNum":"17657806240229","routing":"JFK/
MAD"}
```

Example 3: Select all bag tag numbers originating from SFO/transit through SFO.

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

Explanation: You fetch the tag number of bags that either originate from SFO or pass through SFO. Though you are using the value comparison operator **=**, since the `flightLegs` is an array, the left operand of comparison operator **=** is a sequence with more than one item. That is the reason to use the sequence operator **any** in addition to the value comparison operator **=**. Else you get the following error.

```
Error handling command SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc
FROM BaggageInfo bag WHERE bag.bagInfo[].flightLegs[].fltRouteSrc= "SFO":
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.
```

Output:

```
{"tagNum":"17657806240001","fltRouteSrc":"SFO"}
{"tagNum":"17657806216554","fltRouteSrc":"SFO"}
{"tagNum":["17657806240001","17657806340001"],"fltRouteSrc":["SFO","SFO"]}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"tagNum":"1765780626568","fltRouteSrc":"SFO"}
{"tagNum":"17657806247861","fltRouteSrc":"SFO"}
```

Example 4: Select all bag tag numbers which did not originate from JFK.

```
SELECT bag.bagInfo[].tagNum,
bag.bagInfo[].flightLegs[0].fltRouteSrc
FROM BaggageInfo bag
WHERE bag.bagInfo.flightLegs[0].fltRouteSrc!=ANY "JFK"
```

Explanation: The assumption here is that the first record of the `flightLegs` array has the details of the source location. You fetch the tag number of bags that did not originate from JFK and so using a **!=** operator here. Though you are using the value comparison operator **!=**, since the `flightLegs` is an array, the left operand of the comparison operator **!=** is a sequence with more than one item. That is the reason to use the sequence operator **any** in addition to the value comparison operator **!=**. Else you get the following error.

```
Error handling command SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[0].fltRouteSrc
FROM BaggageInfo bag WHERE bag.bagInfo.flightLegs[0].fltRouteSrc!="JFK":
Failed to display result set: Error: at (2, 0) 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.
```

Output:

```
{"tagNum":"17657806240001","fltRouteSrc":["SFO","IST","ATH"]}
{"tagNum":"17657806224224","fltRouteSrc":["GRU","ORD"]}
{"tagNum":"17657806243578","fltRouteSrc":["MSQ","FRA"]}
{"tagNum":"1765780623244","fltRouteSrc":["MXP","CDG","SLC"]}
{"tagNum":"17657806228676","fltRouteSrc":["MIA","LAX"]}
{"tagNum":"17657806234185","fltRouteSrc":["MEL","LAX"]}
{"tagNum":"17657806255507","fltRouteSrc":["MXP","CDG","SLC"]}
{"tagNum":"17657806292229","fltRouteSrc":["LAX","TPE"]}
{"tagNum":"17657806255823","fltRouteSrc":["MSQ","FRA"]}
{"tagNum":"17657806247861","fltRouteSrc":["SFO","IST","ATH"]}
{"tagNum":"17657806299833","fltRouteSrc":["YYZ","HKG"]}
{"tagNum":"17657806288937","fltRouteSrc":["MIA","LAX"]}
{"tagNum":"17657806216554","fltRouteSrc":["SFO","ORD"]}
{"tagNum":["17657806240001","17657806340001"],"fltRouteSrc":
["SFO","IST","ATH","SFO","IST","ATH"]}
{"tagNum":"1765780626568","fltRouteSrc":["SFO","IST","ATH"]}
{"tagNum":"17657806255240","fltRouteSrc":["MIA","LAX"]}
{"tagNum":"17657806232501","fltRouteSrc":["BZN","SEA","CDG"]}
{"tagNum":"17657806292518","fltRouteSrc":["MIA","LAX"]}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

BETWEEN Operator

The BETWEEN operator is used to check 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 operation returns a TRUE value if both the expressions return TRUE. If either of the expressions is NULL or leads to a NULL value, the result of the operation is also NULL. The operation returns a FALSE value if any one of the expressions returns FALSE. 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.

Example 1: 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":"QB1O0J","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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Clements","RESERVATION":"XT1O7T","ROUTINGINFO":"MXP/CDG/SLC/BZN"}
{"FULLNAME":"Raymond Griffin","RESERVATION":"XT6K7M","ROUTINGINFO":"MSQ/FRA/HKG"}
{"FULLNAME":"Dierdre Amador","RESERVATION":"ZG8Z5N","ROUTINGINFO":"JFK/MAD"}

Example 2: 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

The IN operator is essentially a compact alternative to a number of OR-ed equality conditions. This operator allows you to specify multiple values in a WHERE clause.

Example: Fetch tag number for the customers `"Lucinda Beckman", "Adam Phillips","Zina Christenson","Fallon Clements"`.

```
SELECT bagdet.fullName, bagdet.bagInfo[].tagNum
FROM BaggageInfo bagdet
WHERE bagdet.fullName IN
("Lucinda Beckman", "Adam Phillips","Zina Christenson","Fallon Clements")
```

Explanation: You fetch the tag numbers of a list of passengers. The list of passengers to be fetched can be given inside an **IN** clause.

Output:

{"fullName":"Lucinda Beckman","tagNum":"17657806240001"}
{"fullName":"Zina Christenson","tagNum":"17657806228676"}
{"fullName":"Adam Phillips","tagNum":"17657806255240"}
{"fullName":"Fallon Clements","tagNum":"17657806255507"}

Copyright © 2022, 2025, Oracle and/or its affiliates.

Regular Expression Conditions

A regular expression is a pattern that the regular expression engine attempts to match with an input string. The `regex_like` function performs regular expression matching. 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.

The pattern string is the regular expression against which the input text is matched. The period (.) is a meta-character that matches every character except a new line. The greedy quantifier (*) is a meta-character that indicates zero or more occurrences of the preceding element. For example, the regex "D.*" matches any string that starts with the character 'D' and is followed by zero or more characters.

Example 1: Fetch baggage information of passengers whose names start with 'Z'.

```
SELECT bag.fullname,bag.bagInfo[].tagNum
FROM BaggageInfo bag
WHERE regex_like(fullName, "Z.*")
```

Explanation: You fetch the full name and tag numbers of passengers whose full name starts with Z. You use a regular expression and specify that the first character in the full name should be "**Z**" and the rest can be anything else.

Output:

```
{"fullname":"Zina Christenson","tagNum":"17657806228676"}
{"fullname":"Zulema Martindale","tagNum":"17657806288937"}
```

Example 2: Fetch baggage information of passengers whose flight source location has an "**M**" in it.

Option 1:

```
SELECT bag.fullname,bag.bagInfo[].tagNum,
bag.bagInfo[].flightLegs[0].fltRouteSrc
FROM BaggageInfo bag
WHERE regex_like(bag.bagInfo.flightLegs[0].fltRouteSrc, ".*M.*")
```

Explanation: The assumption here is that the first record of the `flightLegs` array has the details of the source location. You fetch the full name and tag numbers of passengers whose flight source has an "**M**" in it. You use a regular expression and specify that one of the characters in the source field should be "**M**" and the rest can be anything else.

You can also use different approaches to write queries to solve the above problem.

Option 2: Instead of hard coding the index of the `flightLegs` array, you use the `regex_like` function to determine the correct index.

```
SELECT bag.fullname,bag.bagInfo[].tagNum,
bag.bagInfo[].flightLegs[].fltRouteSrc
FROM BaggageInfo bag
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
WHERE EXISTS (bag.bagInfo.flightLegs[regex_like($element.fltRouteSrc,
".*M.*")])
```

Option 3: You use the substring of the "routing" field to extract the source and then use `regex_like` function to search the letter **M** in the source.

```
SELECT bag.fullname,bag.bagInfo[].tagNum,
substring(bag.bagInfo[].routing,0,3)
FROM BaggageInfo bag WHERE
regex_like(substring(bag.bagInfo[].routing,0,3), ".*M.*")
```

Output:

```
{"fullname":"Raymond Griffin","tagNum":"17657806243578","fltRouteSrc":"MSQ"}
{"fullname":"Elane Lemons","tagNum":"1765780623244","fltRouteSrc":"MXP"}
{"fullname":"Zina Christenson","tagNum":"17657806228676","fltRouteSrc":"MIA"}
{"fullname":"Zulema Martindale","tagNum":"17657806288937","fltRouteSrc":"MIA"}
{"fullname":"Adam Phillips","tagNum":"17657806255240","fltRouteSrc":"MIA"}
{"fullname":"Joanne Diaz","tagNum":"17657806292518","fltRouteSrc":"MIA"}
{"fullname":"Teena Colley","tagNum":"17657806255823","fltRouteSrc":"MSQ"}
{"fullname":"Omar Harvey","tagNum":"17657806234185","fltRouteSrc":"MEL"}
{"fullname":"Fallon Clements","tagNum":"17657806255507","fltRouteSrc":"MXP"}
```

EXISTS Operator

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.

Example 1: Select passenger details and baggage information for those passengers who have three flight segments.

```
SELECT fullName, bag.bagInfo[].tagNum,
bag.bagInfo[].routing
FROM BaggageInfo bag
WHERE EXISTS bag.bagInfo[].flightLegs[2]
```

Explanation: You fetch the details of the passengers who have three flight segments. You determine this by evaluating if the third element of the flight legs array is present using the `EXISTS` operator.

Output:

```
{"fullName":"Lorenzo Phil","tagNum":
["17657806240001","17657806340001"],"routing":["SFO/IST/ATH/JTR","SFO/IST/ATH/
JTR"]}
{"fullName":"Gerard Greene","tagNum":"1765780626568","routing":"SFO/IST/ATH/
JTR"}
{"fullName":"Doris Martin","tagNum":"17657806232501","routing":"BZN/SEA/CDG/
MXP"}
{"fullName":"Fallon
Clements","tagNum":"17657806255507","routing":"MXP/CDG/SLC/BZN"}
{"fullName":"Michelle Payne","tagNum":"17657806247861","routing":"SFO/IST/ATH/
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
JTR"}
{"fullName":"Lucinda
Beckman","tagNum":"17657806240001","routing":"SFO/IST/ATH/JTR"}
{"fullName":"Elane Lemons","tagNum":"1765780623244","routing":"MXP/CDG/SLC/
BZN"}
```

Example 2: Fetch the full name and tag number for all customer baggage shipped after 2019.

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

Explanation: The bag arrival date value for every bag should be greater than the year 2019. Here the **"$element"** is bound to the context row (every bag of the customer). The `EXISTS` operator checks whether the sequence returned by its input expression is empty or not. The sequence returned by the comparison operator ">=" is non-empty for all bags which arrived after 2019.

Output:

```
{"fullName":"Lucinda Beckman","ticketNo":1762320569757}
{"fullName":"Adelaide Willard","ticketNo":1762392135540}
{"fullName":"Raymond Griffin","ticketNo":1762399766476}
{"fullName":"Elane Lemons","ticketNo":1762324912391}
{"fullName":"Zina Christenson","ticketNo":1762390789239}
{"fullName":"Zulema Martindale","ticketNo":1762340579411}
{"fullName":"Dierdre Amador","ticketNo":1762376407826}
{"fullName":"Henry Jenkins","ticketNo":176234463813}
{"fullName":"Rosalia Triplett","ticketNo":1762311547917}
{"fullName":"Lorenzo Phil","ticketNo":1762320369957}
{"fullName":"Gerard Greene","ticketNo":1762341772625}
{"fullName":"Adam Phillips","ticketNo":1762344493810}
{"fullName":"Doris Martin","ticketNo":1762355527825}
{"fullName":"Joanne Diaz","ticketNo":1762383911861}
{"fullName":"Omar Harvey","ticketNo":1762348904343}
{"fullName":"Fallon Clements","ticketNo":1762350390409}
{"fullName":"Lisbeth Wampler","ticketNo":1762355854464}
{"fullName":"Teena Colley","ticketNo":1762357254392}
{"fullName":"Michelle Payne","ticketNo":1762330498104}
{"fullName":"Mary Watson","ticketNo":1762340683564}
{"fullName":"Kendal Biddle","ticketNo":1762377974281}
```

Is-Of-Type Operator

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.

Example: Fetch the names of the passengers whose baggage tags contain only numbers and not a STRING.

```
SELECT fullname,bag.bagInfo.tagNum
FROM  BaggageInfo bag
WHERE bag.bagInfo.tagNum is of type (NUMBER)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Explanation: The `tagNum` in the `bagInfo` schema is a STRING data type. But the application could take in a NUMBER value as `tagNum` by mistake. The query captures the passengers for whom the `tagNum` column has only numbers.

Output:

```
{"fullname":"Raymond Griffin","tagNum":17657806243578}
```

If you query the `bagInfo` schema for the above `tagNum` as STRING, no rows are displayed.

```
SELECT * FROM BaggageInfo bag WHERE tagnum = "17657806232501"
0 row returned
```

You can also fetch the names of the passengers whose baggage tags contain only STRING.

```
SELECT fullname,bag.bagInfo.tagNum
FROM BaggageInfo bag
WHERE  bag.bagInfo.tagNum is of type (STRING)
```

SQL Operators examples using QueryRequest API

You can use `QueryRequest` API and filter data from a NoSQL table using SQL operators.

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***SQLOperators.java*** from the examples here.

```
 //Fetch rows from the table
private static void fetchRows(NoSQLHandle handle,String sqlstmt) throws
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)){
      for (MapValue res : results) {
         System.out.println("\t" + res);
      }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}



String seq_comp_ope="SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc FROM BaggageInfo
bag WHERE bag.bagInfo[].flightLegs[].fltRouteSrc=any \"SFO\"";
System.out.println("Using Sequence Comparison operator ");
fetchRows(handle,seq_comp_ope);
String logical_ope="SELECT fullName, bag.bagInfo[].tagNum,
bag.bagInfo[].routing,bag.bagInfo[].flightLegs[].fltRouteSrc FROM BaggageInfo
bag WHERE NOT bag.bagInfo[].flightLegs[].fltRouteSrc=any \"SFO\"";
System.out.println("Using Logical operator ");
fetchRows(handle,logical_ope);
String value_comp_ope="SELECT fullname, bag.bagInfo[].routing FROM
BaggageInfo bag WHERE gender=\"M\"";
System.out.println("Using Value Comparison operator ");
fetchRows(handle,value_comp_ope);
String in_ope="SELECT bagdet.fullName, bagdet.bagInfo[].tagNum FROM
BaggageInfo bagdet WHERE bagdet.fullName IN (\"Lucinda Beckman\", \"Adam
Phillips\",\"Dierdre Amador\",\"Fallon Clements\")";System.out.println("Using
IN operator ");fetchRows(handle,in_ope);
String exists_ope="SELECT fullName, bag.ticketNo FROM BaggageInfo bag WHERE
EXISTS bag.bagInfo[$element.bagArrivalDate >=\"2019-03-01T00:00:00\"]";
System.out.println("Using EXISTS operator ");
fetchRows(handle,exists_ope);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***SQLOperators.py*** from the examples here.

```
# Fetch data from the table
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))


seqcomp_stmt = '''SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc
                     FROM BaggageInfo bag WHERE
bag.bagInfo[].flightLegs[].fltRouteSrc=any "SFO"'''
print('Using Sequence Comparison operator:')
fetch_data(handle,seqcomp_stmt)
logope_stmt = '''SELECT fullName, bag.bagInfo[].tagNum, bag.bagInfo[].routing,
                 bag.bagInfo[].flightLegs[].fltRouteSrc
                 FROM BaggageInfo bag
                 WHERE NOT bag.bagInfo[].flightLegs[].fltRouteSrc=any "SFO"'''
print('Using Logical operator:')
fetch_data(handle,logope_stmt)
valcomp_stmt = '''SELECT fullname, bag.bagInfo[].routing
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                  FROM BaggageInfo bag WHERE gender="M"'''
print('Using Value Comparison operator:')
fetch_data(handle,valcomp_stmt)
inope_stmt = '''SELECT bagdet.fullName, bagdet.bagInfo[].tagNum
                FROM BaggageInfo bagdet WHERE bagdet.fullName IN
               ("Lucinda Beckman", "Adam Phillips","Dierdre Amador","Fallon
Clements")'''
print('Using IN operator:')
fetch_data(handle,inope_stmt)
existsope_stmt = '''SELECT fullName, bag.ticketNo FROM BaggageInfo bag WHERE
                    EXISTS bag.bagInfo[$element.bagArrivalDate
>="2019-03-01T00:00:00"]'''
print('Using EXISTS operator:')
fetch_data(handle,existsope_stmt)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***SQLOperators.go*** from the examples here.

```
 //fetch data from the table
func fetchData(client *nosqldb.Client, err error, tableName string, querystmt
string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,    }
   var results []*types.MapValue
   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()
      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
      }
      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
   }
   for i, r := range results {
      fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}


seqcomp_stmt := `SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc
                    FROM BaggageInfo bag WHERE
bag.bagInfo[].flightLegs[].fltRouteSrc=any "SFO"`
fmt.Printf("Using Sequence Comparison operator:\n"
fetchData(client, err,tableName,seqcomp_stmt)

logope_stmt := `SELECT fullName, bag.bagInfo[].tagNum, bag.bagInfo[].routing,
                bag.bagInfo[].flightLegs[].fltRouteSrc
                FROM BaggageInfo bag
                WHERE NOT bag.bagInfo[].flightLegs[].fltRouteSrc=any "SFO"`
fmt.Printf("Using Logical operator:\n")
fetchData(client, err,tableName,logope_stmt)

valcomp_stmt := `SELECT fullname, bag.bagInfo[].routing FROM BaggageInfo bag
WHERE gender="M"`
fmt.Printf("Using Value Comparison operator:\n")
fetchData(client, err,tableName,valcomp_stmt)

inope_stmt := `SELECT bagdet.fullName, bagdet.bagInfo[].tagNum FROM
BaggageInfo bagdet
            WHERE bagdet.fullName IN ("Lucinda Beckman", "Adam
Phillips","Dierdre Amador","Fallon Clements")`
fmt.Printf("Using IN operator:\n")
fetchData(client, err,tableName,inope_stmt)

existsope_stmt := `SELECT fullName, bag.ticketNo FROM BaggageInfo bag WHERE
                     EXISTS bag.bagInfo[$element.bagArrivalDate
>="2019-03-01T00:00:00"]`
fmt.Printf("Using EXISTS operator:\n")
fetchData(client, err,tableName,existsope_stmt)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***SQLOperators.js*** from the examples here.

```
  //fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}
```

TypeScript: Download the full code ***SQLOperators.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


const seqcomp_stmt = `SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc
                      FROM BaggageInfo bag WHERE
bag.bagInfo[].flightLegs[].fltRouteSrc=any "SFO"`
console.log("Using Sequence Comparison operator");
await fetchData(handle,seqcomp_stmt);

const logope_stmt = `SELECT fullName, bag.bagInfo[].tagNum,
bag.bagInfo[].routing,
                     bag.bagInfo[].flightLegs[].fltRouteSrc
                     FROM BaggageInfo bag
                     WHERE NOT bag.bagInfo[].flightLegs[].fltRouteSrc=any
"SFO"`
console.log("Using Logical operator");
await fetchData(handle,logope_stmt);

const valcomp_stmt = `SELECT fullname, bag.bagInfo[].routing FROM BaggageInfo
bag WHERE gender="M"`
console.log("Using Value Comparison operator");
await fetchData(handle,valcomp_stmt);

const inope_stmt = `SELECT bagdet.fullName, bagdet.bagInfo[].tagNum
                    FROM BaggageInfo bagdet WHERE bagdet.fullName IN
                    ("Lucinda Beckman", "Adam Phillips","Dierdre
Amador","Fallon Clements")`
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
console.log("Using IN operator");
await fetchData(handle,inope_stmt);

const existsope_stmt = `SELECT fullName, bag.ticketNo FROM BaggageInfo bag
WHERE
                        EXISTS bag.bagInfo[$element.bagArrivalDate
>="2019-03-01T00:00:00"]`
console.log("Using EXISTS operator");
await fetchData(handle,existsope_stmt);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***SQLOperators.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}


private const string seqcomp_stmt =@"SELECT
bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc
                                           FROM BaggageInfo bag WHERE
bag.bagInfo[].flightLegs[].fltRouteSrc=any ""SFO""";
Console.WriteLine("\nUsing Sequence Comparison operator!");
await fetchData(client,seqcomp_stmt);

private const string logope_stmt =@"SELECT fullName, bag.bagInfo[].tagNum,
bag.bagInfo[].routing,

bag.bagInfo[].flightLegs[].fltRouteSrc
                                           FROM BaggageInfo bag
                                           WHERE NOT
bag.bagInfo[].flightLegs[].fltRouteSrc=any ""SFO""";
Console.WriteLine("\nUsing Logical operator!");
await fetchData(client,logope_stmt);

private const string valcomp_stmt =@"SELECT fullname, bag.bagInfo[].routing
FROM BaggageInfo bag WHERE gender=""M""" ;
Console.WriteLine("\nUsing Value Comparison operator!");
await fetchData(client,valcomp_stmt);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
private const string inope_stmt =@"SELECT bagdet.fullName,
bagdet.bagInfo[].tagNum
                                    FROM BaggageInfo bagdet WHERE
bagdet.fullName IN
                                    (""Lucinda Beckman"", ""Adam
Phillips"",""Dierdre Amador"",""Fallon Clements"")";
Console.WriteLine("\nUsing IN operator!");
await fetchData(client,inope_stmt);

private const string existsope_stmt =@"SELECT fullName, bag.ticketNo FROM
BaggageInfo bag WHERE
                                       EXISTS
bag.bagInfo[$element.bagArrivalDate >=""2019-03-01T00:00:00""]";
Console.WriteLine("\nUsing EXISTS operator!");
await fetchData(client,existsope_stmt);
```

Sorting, Grouping & Limiting results

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Ordering results
- Limit and offset results
- Grouping results
- Aggregating results
- Examples using QueryRequest API

Ordering results

Use the ORDER BY clause to order the results by any column, primary key or non-primary key.

Example 1: Sort the ticket number of all passengers by their full name.

```
SELECT bag.ticketNo, bag.fullName
FROM BaggageInfo bag
ORDER BY bag.fullName
```

Explanation: You are sorting the ticket number of passengers in the `BaggageInfo` schema based on the full name of the passengers in ascending order.

Output:

```
{"ticketNo":1762344493810,"fullName":"Adam Phillips"}
{"ticketNo":1762392135540,"fullName":"Adelaide Willard"}
{"ticketNo":1762376407826,"fullName":"Dierdre Amador"}
{"ticketNo":1762355527825,"fullName":"Doris Martin"}
{"ticketNo":1762324912391,"fullName":"Elane Lemons"}
{"ticketNo":1762350390409,"fullName":"Fallon Clements"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{"ticketNo":1762341772625,"fullName":"Gerard Greene"}
{"ticketNo":176234463813,"fullName":"Henry Jenkins"}
{"ticketNo":1762383911861,"fullName":"Joanne Diaz"}
{"ticketNo":1762377974281,"fullName":"Kendal Biddle"}
{"ticketNo":1762355854464,"fullName":"Lisbeth Wampler"}
{"ticketNo":1762320369957,"fullName":"Lorenzo Phil"}
{"ticketNo":1762320569757,"fullName":"Lucinda Beckman"}
{"ticketNo":1762340683564,"fullName":"Mary Watson"}
{"ticketNo":1762330498104,"fullName":"Michelle Payne"}
{"ticketNo":1762348904343,"fullName":"Omar Harvey"}
{"ticketNo":1762399766476,"fullName":"Raymond Griffin"}
{"ticketNo":1762311547917,"fullName":"Rosalia Triplett"}
{"ticketNo":1762357254392,"fullName":"Teena Colley"}
{"ticketNo":1762390789239,"fullName":"Zina Christenson"}
{"ticketNo":1762340579411,"fullName":"Zulema Martindale"}
```

Example 2: Fetch the passenger details(full name, tag number) by the last seen time (latest first) for passengers (sorted by their name) whose last seen station is **MEL**.

```
SELECT bag.fullName, bag.bagInfo[].tagNum,
bag.bagInfo[].lastSeenTimeGmt
FROM BaggageInfo bag
WHERE bag.bagInfo[].lastSeenStation=any "MEL"
ORDER BY bag.bagInfo[].lastSeenTimeGmt DESC
```

Explanation: You first filter the data in the `BaggageInfo` table based on the last seen station and you sort the filtered results based on the last seen time and the full name of the passengers in descending order. You do this using the ORDER BY clause.

Note

You can use more than one column to sort the output of the query.

Output:

```
{"fullName":"Adam
Phillips","tagNum":"17657806255240","lastSeenTimeGmt":"2019-02-01T16:13:00Z"}
{"fullName":"Zina
Christenson","tagNum":"17657806228676","lastSeenTimeGmt":"2019-02-04T10:08:00Z
"}
{"fullName":"Joanne
Diaz","tagNum":"17657806292518","lastSeenTimeGmt":"2019-02-16T16:13:00Z"}
{"fullName":"Zulema
Martindale","tagNum":"17657806288937","lastSeenTimeGmt":"2019-02-25T20:15:00Z"
}
```

Limit and offset results

Use the `LIMIT` clause to limit the number of results returned from a SELECT statement. For example, if there are 1000 rows in a table, limit the number of rows to return by specifying a LIMIT value. It is recommended to use LIMIT and OFFSET with an ORDER BY clause. Otherwise, the results are returned in a random order, producing unpredictable results.

Copyright © 2022, 2025, Oracle and/or its affiliates.

A good use-case/example of using LIMIT and OFFSET is the application paging of results. Say for example your application wants to show 4 results per page. You can use limit and offset to implement stateless paging in the application. If you are showing n (say 4) results per page, then the results for page m (say 2) are being displayed, then offset would be (n*m-1) which is 4 in this example and the limit would be n(which is 4 here).

Example 1: Your application can show 4 results on a page. Fetch the details fetched by your application in the first page for passengers whose last seen station is **JTR**.

```
SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time
FROM BaggageInfo $bag,
$bag.bagInfo[].lastSeenTimeGmt $flt_time
WHERE $bag.bagInfo[].lastSeenStation=any "JTR"
ORDER BY $flt_time LIMIT 4
```

Explanation: You filter the data in the `BaggageInfo` table based on the last seen station and you sort the result based on the last seen time. You use an unnest array to flatten your data. That is the `bagInfo` array is flattened and the last seen time is fetched. You need to just display the first 4 rows from the result set.

Output:

```
{"fullName":"Michelle
Payne","tagNum":"17657806247861","flt_time":"2019-02-02T23:59:00Z"}
{"fullName":"Gerard
Greene","tagNum":"1765780626568","flt_time":"2019-03-07T16:01:00Z"}
{"fullName":"Lorenzo Phil","tagNum":
["17657806240001","17657806340001"],"flt_time":"2019-03-12T15:05:00Z"}
{"fullName":"Lucinda
Beckman","tagNum":"17657806240001","flt_time":"2019-03-12T15:05:00Z"}
```

Example 2: Your application can show 4 results on a page. Fetch the details fetched by your application in the second page for passengers whose last seen station is **JTR**.

```
SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time
FROM BaggageInfo $bag,
$bag.bagInfo[].lastSeenTimeGmt $flt_time
WHERE $bag.bagInfo[].lastSeenStation=any "JTR"
ORDER BY $flt_time LIMIT 4 OFFSET 4
```

Explanation: You filter the data in the `BaggageInfo` table based on the last seen station and you sort the result based on the last seen time. You use an unnest array to flatten your data. You need to display the contents of the second page, so you set an OFFSET 4. Though you LIMIT to 4 rows, only one row is displayed as the total result set is only 5. The first few are skipped and the fifth one is displayed.

Output:

```
{"fullName":"Lorenzo Phil","tagNum":["17657806240001","17657806340001"],
"flt_time":"2019-03-12T16:05:00Z"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Grouping results

Use the GROUP BY clause to group the results by one or more table columns. Typically, a GROUP BY clause is used in conjunction with an aggregate expression such as COUNT, SUM, and AVG.

Example 1: Display the number of bags for each reservation made.

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

Explanation: Every passenger has one reservation code (`confNo`). A passenger can have more than one baggage. Here you group the data based on the reservation code and you get the count of the `bagInfo` array which gives the number of bags per reservation.

Output:

```
{"confNo":"FH7G1W","TOTAL_BAGS":1}
{"confNo":"PQ1M8N","TOTAL_BAGS":1}
{"confNo":"XT6K7M","TOTAL_BAGS":1}
{"confNo":"DN3I4Q","TOTAL_BAGS":1}
{"confNo":"QB1O0J","TOTAL_BAGS":1}
{"confNo":"TX1P7E","TOTAL_BAGS":1}
{"confNo":"CG6O1M","TOTAL_BAGS":1}
{"confNo":"OH2F8U","TOTAL_BAGS":1}
{"confNo":"BO5G3H","TOTAL_BAGS":1}
{"confNo":"ZG8Z5N","TOTAL_BAGS":1}
{"confNo":"LE6J4Z","TOTAL_BAGS":1}
{"confNo":"XT1O7T","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 2: Select the total baggage originating from each airport (excluding the transit baggage).

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

Explanation: You want to 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. So you group the data with the flight source values of the first record of the `flightLegs` array(as the first record is the source). You then determine the count of baggage.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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}
```

Aggregating results

Use the built in aggregate and sequence aggregate functions to find information such as a count, a sum, an average, a minimum, or a maximum.

Example 1: 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 expression '=' 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}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 2: 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"}
```

Examples using QueryRequest API

You can use `QueryRequest` API to group and order data and also fetch it from a NoSQL table.

- Java
- Python
- Go
- Node.js

Copyright © 2022, 2025, Oracle and/or its affiliates.

- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***GroupSortData.java*** from the examples here.

```
 //Fetch rows from the table
private static void fetchRows(NoSQLHandle handle,String sqlstmt) throws
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)){
      for (MapValue res : results) {
         System.out.println("\t" + res);
      }
   }
}


String orderby_stmt="SELECT bag.fullName,
bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenTimeGmt FROM BaggageInfo bag "+
                             "WHERE bag.bagInfo[].lastSeenStation=any \"MEL\"
ORDER BY bag.bagInfo[].lastSeenTimeGmt DESC";
System.out.println("Using ORDER BY to sort data ");
fetchRows(handle,orderby_stmt);
String sortlimit_stmt="SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time
FROM BaggageInfo $bag, "+
                       "$bag.bagInfo[].lastSeenTimeGmt $flt_time
WHERE $bag.bagInfo[].lastSeenStation=any \"JTR\""+
                       "ORDER BY $flt_time LIMIT 4";
System.out.println("Using ORDER BY and LIMIT to sort and limit data ");
fetchRows(handle,sortlimit_stmt);
String groupsortlimit_stmt="SELECT $flt_src as SOURCE,count(*) as COUNT FROM
BaggageInfo $bag, "+
                           "$bag.bagInfo.flightLegs[0].fltRouteSrc $flt_src
GROUP BY $flt_src";
System.out.println("Using GROUP BY, ORDER BY and LIMIT to group, sort and
limit data ");
fetchRows(handle,groupsortlimit_stmt);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***GroupSortData.py*** from the examples here.

```
# Fetch data from the table
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   for r in result.get_results():
      print('\t' + str(r))


orderby_stmt = '''SELECT bag.fullName,
bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenTimeGmt FROM BaggageInfo bag
                     WHERE bag.bagInfo[].lastSeenStation=any \"MEL\" ORDER BY
bag.bagInfo[].lastSeenTimeGmt DESC'''
print('Using ORDER BY to sort data:')
fetch_data(handle,orderby_stmt)

sortlimit_stmt = '''SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time FROM
BaggageInfo $bag,
                    $bag.bagInfo[].lastSeenTimeGmt $flt_time
                    WHERE $bag.bagInfo[].lastSeenStation=any "JTR"
                    ORDER BY $flt_time LIMIT 4'''
print('Using ORDER BY and LIMIT to sort and limit data:')
fetch_data(handle,sortlimit_stmt)

groupsortlimit_stmt = '''SELECT $flt_src as SOURCE, count(*) as COUNT FROM
BaggageInfo $bag,
                         $bag.bagInfo.flightLegs[0].fltRouteSrc $flt_src
GROUP BY $flt_src'''
print('Using GROUP BY, ORDER BY and LIMIT to group, sort and limit data:')
fetch_data(handle,groupsortlimit_stmt)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***GroupSortData.go*** from the examples here.

```
 //fetch data from the table
func fetchData(client *nosqldb.Client, err error, tableName string, querystmt
string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,    }
   var results []*types.MapValue
   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()
      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        }
        results = append(results, res...)
        if queryReq.IsDone() {
            break
        }
    }
    for i, r := range results {
        fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
    }
}


orderby_stmt := `SELECT bag.fullName,
bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenTimeGmt FROM BaggageInfo bag
                    WHERE bag.bagInfo[].lastSeenStation=any "MEL" ORDER BY
bag.bagInfo[].lastSeenTimeGmt DESC`
fmt.Printf("Using ORDER BY to sort data::\n")
fetchData(client, err,tableName,orderby_stmt)

sortlimit_stmt := `SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time FROM
BaggageInfo $bag,
                   $bag.bagInfo[].lastSeenTimeGmt $flt_time
                   WHERE $bag.bagInfo[].lastSeenStation=any "JTR"
                   ORDER BY $flt_time LIMIT 4`
fmt.Printf("Using ORDER BY and LIMIT to sort and limit data::\n")
fetchData(client, err,tableName,sortlimit_stmt)

groupsortlimit_stmt := `SELECT $flt_src as SOURCE, count(*) as COUNT FROM
BaggageInfo $bag,
                        $bag.bagInfo.flightLegs[0].fltRouteSrc $flt_src GROUP
BY $flt_src`
fmt.Printf("Using GROUP BY, ORDER BY and LIMIT to group, sort and limit
data::\n")
fetchData(client, err,tableName,groupsortlimit_stmt)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***GroupSortData.js*** from the examples here.

```
  //fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}
```

TypeScript: Download the full code ***GroupSortData.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


const orderby_stmt = `SELECT bag.fullName,
bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenTimeGmt FROM BaggageInfo bag
                      WHERE bag.bagInfo[].lastSeenStation=any \"MEL\" ORDER
BY bag.bagInfo[].lastSeenTimeGmt DESC`
console.log("Using ORDER BY to sort data");
await fetchData(handle,orderby_stmt);

const sortlimit_stmt = `SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time
FROM BaggageInfo $bag,
                        $bag.bagInfo[].lastSeenTimeGmt $flt_time
                        WHERE $bag.bagInfo[].lastSeenStation=any "JTR"
                        ORDER BY $flt_time LIMIT 4`
console.log("Using ORDER BY and LIMIT to sort and limit data");
await fetchData(handle,sortlimit_stmt);

const groupsortlimit_stmt = `SELECT $flt_src as SOURCE, count(*) as COUNT
FROM BaggageInfo $bag,
                             $bag.bagInfo.flightLegs[0].fltRouteSrc $flt_src
GROUP BY $flt_src`
console.log("Using GROUP BY, ORDER BY and LIMIT to group, sort and limit
data");
await fetchData(handle,groupsortlimit_stmt);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***GroupSortData.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}


private const string orderby_stmt =@"SELECT bag.fullName,
bag.bagInfo[].tagNum,bag.bagInfo[].lastSeenTimeGmt
                                             FROM BaggageInfo bag WHERE
bag.bagInfo[].lastSeenStation=any ""MEL""
                                             ORDER BY
bag.bagInfo[].lastSeenTimeGmt DESC";
Console.WriteLine("\nUsing ORDER BY to sort data!");
await fetchData(client,orderby_stmt);

private const string sortlimit_stmt
=@"SELECT $bag.fullName, $bag.bagInfo.tagNum, $flt_time FROM BaggageInfo $bag,
                                             $bag.bagInfo[].lastSeenTimeGmt $f
lt_time

WHERE $bag.bagInfo[].lastSeenStation=any ""JTR""
                                             ORDER BY $flt_time LIMIT 4";
Console.WriteLine("\nUsing ORDER BY and LIMIT to sort and limit data!");
await fetchData(client,sortlimit_stmt);


private const string groupsortlimit_stmt =@"SELECT $flt_src as SOURCE,
count(*) as COUNT FROM BaggageInfo $bag,
                                                  $bag.bagInfo.flightLegs[0].f
ltRouteSrc $flt_src GROUP BY $flt_src" ;
Console.WriteLine("\nUsing GROUP BY, ORDER BY and LIMIT to group, sort and
limit data:");
await fetchData(client,groupsortlimit_stmt);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Primary Expressions in SQL

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Parenthesized Expressions
- Case Expressions
- Cast Expression
- Sequence Transform Expressions
- Extract Expressions
- SQL Expression examples using QueryRequest API

Parenthesized Expressions

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.

Example: Fetch the full name, tag number, and routing details of passengers either boarding at JFK /traversing through JFK and their destination is either MAD or VIE.

```
SELECT fullName, bag.bagInfo.tagNum,
bag.bagInfo.routing,
bag.bagInfo[].flightLegs[].fltRouteDest
FROM BaggageInfo bag
WHERE bag.bagInfo.flightLegs[].fltRouteSrc=any "JFK" AND
(bag.bagInfo[].flightLegs[].fltRouteDest=any "MAD" OR
bag.bagInfo[].flightLegs[].fltRouteDest=any "VIE" )
```

Explanation: You want to fetch the full name, tag number, and routing details of passengers. The first filter condition is that the boarding point/transit is JFK. Once this is satisfied the second filter condition is that destination is either MAD or VIE. You use an OR condition to filter the destination value.

Output:

```
{"fullName":"Dierdre Amador","tagNum":"17657806240229","routing":"JFK/
MAD","fltRouteDest":"MAD"}
{"fullName":"Rosalia Triplett","tagNum":"17657806215913","routing":"JFK/IST/
VIE","fltRouteDest":["IST","VIE"]}
{"fullName":"Kendal Biddle","tagNum":"17657806296887","routing":"JFK/IST/
VIE","fltRouteDest":["IST","VIE"]}
```

Case Expressions

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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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:

```
SELECT
    fullName,
    CASE
        WHEN NOT exists bag.bagInfo.flightLegs[0]
        THEN "you have no bag info"
        WHEN NOT exists bag.bagInfo.flightLegs[1]
        THEN "you have one hop"
        WHEN NOT exists bag.bagInfo.flightLegs[2]
        THEN "you have two hops."
        ELSE "you have three hops."
    END AS NUMBER_HOPS
FROM BaggageInfo bag WHERE ticketNo=1762340683564
```

Explanation: You want to determine how many transits are there for the passenger `bagInfo` using a CASE statement. If the `flightLegs` array has no elements, then the passenger has no bag data. When the `flightLegs` array has only one element, then there is only one transit point. Similarly, if the `flightLegs` array has two elements, then there is two hops. Else there is three transit points. Here you assume that a bag can have at the most three transit points/ hops.

Output:

```
{"fullName":"Mary Watson","NUMBER_HOPS":"you have two hops."}
```

Example 2: Write a query to alert the system to update the `tagNum` of passengers if the existing value is not a string.

```
SELECT bag.bagInfo[].tagNum,
CASE
   WHEN bag.bagInfo[0].tagNum is of type (NUMBER)
   THEN "Tagnumber is not a STRING. Update the data"
   ELSE "Tagnumber has correct datatype"
   END AS tag_NUM_TYPE
FROM BaggageInfo bag
```

Explanation: The `tagNum` of passengers in the `bagInfo` schema is a STRING data type. But the application could take in a NUMBER value as the value of `tagNum` by mistake. The query uses "is of type" operator to capture this and prompts the system to update the `tagNum` if the existing value is not a string.

Output (only few rows are shown for brevity).

```
{"tagNum":"17657806240001","tag_NUM_TYPE":"Tagnumber has correct datatype"}
{"tagNum":"17657806224224","tag_NUM_TYPE":"Tagnumber has correct datatype"}
{"tagNum":17657806243578,"tag_NUM_TYPE":"Tagnumber is not a STRING. Update
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
the data"}
{"tagNum":"1765780623244","tag_NUM_TYPE":"Tagnumber has correct datatype"}
```

Cast Expression

The cast expression creates, if possible, new items of a given target type from the items of its input sequence. For example, a STRING can be converted to TIMESTAMP(0) using CAST expression.

Rules followed in a CAST expression:

- 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 wild card type, the cast is a no-op.
- If the target type is JSON, then an error is raised if the input item is a non-json atomic type.
- If the target type is an array type, an error is raised if the input item is not an array.
- If the target type is string, the input item may be of any type. That means every item can be cast to a string. For timestamps, their string value is in UTC and has the format `uuuu-MM-dd['T'HH:mm:ss]`.
- If the target type is an atomic type other than string, the input item must also be atomic.
  - * 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 that the input item, the resulting timestamp is the one closest to the input timestamp in the target precision.
- To cast a STRING to TIMESTAMP, if the input has STRING values in ISO-8601 format, then it will be automatically converted by the SQL runtime into TIMESTAMP data type.

Copyright © 2022, 2025, Oracle and/or its affiliates.

> **Note**
>
> ISO8601 describes an internationally accepted way to represent dates, times, and durations.
> **Syntax:** Date with time: YYYY-MM-DDThh:mm:ss[.s[s[s[s[s[s]]]]][Z|(+|-)hh:mm]
>
> where
>
> – YYYY specifies the year, as four decimal digits
> – MM specifies the month, as two decimal digits, 00 to 12
> – DD specifies the day, as two decimal digits, 00 to 31
> – hh specifies the hour, as two decimal digits, 00 to 23
> – mm specifies the minutes, as two decimal digits, 00 to 59
> – ss[.s[s[s[s[s]]]]] specifies the seconds, as two decimal digits, 00 to 59, optionally followed by a decimal point and 1 to 6 decimal digits (representing the fractional part of a second).
> – Z specifies UTC time (time zone 0). (It can also be specified by +00:00, but not by –00:00.)
> – (+|-)hh:mm specifies the time-zone as difference from UTC. (One of + or – is required.)

Example 1: Fetch the bag arrival date for the passenger with a reservation code **DN3I4Q** in TIMESTAMP(3) format.

```
SELECT CAST (bag.bagInfo.bagArrivalDate AS Timestamp(3))
AS BAG_ARRIVING_DATE
FROM BaggageInfo bag WHERE bag.confNo=DN3I4Q
```

Explanation: The `bagArrivalDate` is a STRING. Using CAST you are converting this field into a TIMESTAMP format.

Output:

```
{"BAG_ARRIVING_DATE":"2019-02-15T21:21:00.000Z"}
```

Example 2: Fetch the full name and tag number for all customer baggage shipped after 2019.

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

Explanation: You want to filter and display details of the baggage that are shipped after 2019. The bag arrival date for every element in the `flightLegs` array is compared with the given timestamp (`2019-01-01T00:00:00`). Here the casting is implicit as `bagArrivalDate` is a STRING and is directly compared with a static Timestamp value. An explicit CAST function is not needed when an implicit casting can be done. However, your data should be in the format `YYYY-MM-DDTHH:MI:SS` . You then use the EXISTS condition to check if the bagInfo is present for this timestamp condition.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"fullName":"Kendal
Biddle","ticketNo":1762377974281,"bagArrivalDate":"2019-03-05T12:00:00Z"}
{"fullName":"Lucinda
Beckman","ticketNo":1762320569757,"bagArrivalDate":"2019-03-12T15:05:00Z"}
{"fullName":"Adelaide
Willard","ticketNo":1762392135540,"bagArrivalDate":"2019-02-15T21:21:00Z"}
{"fullName":"Raymond
Griffin","ticketNo":1762399766476,"bagArrivalDate":"2019-02-03T08:09:00Z"}
{"fullName":"Elane
Lemons","ticketNo":1762324912391,"bagArrivalDate":"2019-03-15T10:13:00Z"}
{"fullName":"Zina
Christenson","ticketNo":1762390789239,"bagArrivalDate":"2019-02-04T10:08:00Z"}
{"fullName":"Zulema
Martindale","ticketNo":1762340579411,"bagArrivalDate":"2019-02-25T20:15:00Z"}
{"fullName":"Dierdre
Amador","ticketNo":1762376407826,"bagArrivalDate":"2019-03-07T13:51:00Z"}
{"fullName":"Henry
Jenkins","ticketNo":176234463813,"bagArrivalDate":"2019-03-02T13:18:00Z"}
{"fullName":"Rosalia
Triplett","ticketNo":1762311547917,"bagArrivalDate":"2019-02-12T07:04:00Z"}
{"fullName":"Lorenzo Phil","ticketNo":1762320369957,"bagArrivalDate":
["2019-03-12T15:05:00Z","2019-03-12T16:25:00Z"]}
{"fullName":"Gerard
Greene","ticketNo":1762341772625,"bagArrivalDate":"2019-03-07T16:01:00Z"}
{"fullName":"Adam
Phillips","ticketNo":1762344493810,"bagArrivalDate":"2019-02-01T16:13:00Z"}
{"fullName":"Doris
Martin","ticketNo":1762355527825,"bagArrivalDate":"2019-03-22T10:17:00Z"}
{"fullName":"Joanne
Diaz","ticketNo":1762383911861,"bagArrivalDate":"2019-02-16T16:13:00Z"}
{"fullName":"Teena
Colley","ticketNo":1762357254392,"bagArrivalDate":"2019-02-13T11:15:00Z"}
{"fullName":"Michelle
Payne","ticketNo":1762330498104,"bagArrivalDate":"2019-02-02T23:59:00Z"}
{"fullName":"Mary
Watson","ticketNo":1762340683564,"bagArrivalDate":"2019-03-14T06:22:00Z"}
{"fullName":"Omar
Harvey","ticketNo":1762348904343,"bagArrivalDate":"2019-03-02T16:09:00Z"}
{"fullName":"Fallon
Clements","ticketNo":1762350390409,"bagArrivalDate":"2019-02-21T14:08:00Z"}
{"fullName":"Lisbeth
Wampler","ticketNo":1762355854464,"bagArrivalDate":"2019-02-10T10:01:00Z"}
```

Sequence Transform Expressions

A sequence transform expression transforms a sequence into 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Example: For each `ticketNo`, fetch a flat array containing all the actions performed on the luggage of that `ticketNo`.

```
SELECT seq_transform(l.bagInfo[],
      seq_transform(
          $sq1.flightLegs[],
          seq_transform(
              $sq2.actions[],
              {
                "at" : $sq3.actionAt,
                "action" : $sq3.actionCode,
                "flightNo" : $sq2.flightNo,
                "tagNum" : $sq1.tagNum
              }
          )
      )
  ) AS actions
FROM baggageInfo l WHERE ticketNo=1762340683564
```

Explanation: 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`.

Output:

```
{
 "actions":[
   {"action":"ONLOAD to
HKG","at":"YYZ","flightNo":"BM267","tagNum":"17657806299833"},
   {"action":"BagTag Scan at
YYZ","at":"YYZ","flightNo":"BM267","tagNum":"17657806299833"},
   {"action":"Checkin at
YYZ","at":"YYZ","flightNo":"BM267","tagNum":"17657806299833"},
   {"action":"Offload to Carousel at
BLR","at":"BLR","flightNo":"BM115","tagNum":"17657806299833"},
   {"action":"ONLOAD to
BLR","at":"HKG","flightNo":"BM115","tagNum":"17657806299833"},
   {"action":"OFFLOAD from
HKG","at":"HKG","flightNo":"BM115","tagNum":"17657806299833"}
 ]
}
```

Extract Expressions

The EXTRACT expression extracts a component from a timestamp.

You can specify one of the following keywords to extract the corresponding date part from the timestamp: YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, MILLISECOND, MICROSECOND, NANOSECOND, WEEK, ISOWEEK.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 1: What is the full name and baggage arrival year for the customer with ticket number **1762383911861**.

```
SELECT fullName,
EXTRACT (YEAR FROM CAST (bag.bagInfo.bagArrivalDate AS Timestamp(0)))
AS YEAR FROM BaggageInfo bag
WHERE ticketNo=1762383911861
```

Explanation: You first use CAST to convert the `bagArrivalDate` to a TIMESTAMP and then fetch the YEAR component from the timestamp.

Output:

```
{"fullName":"Joanne Diaz","YEAR":2019}
```

Example 2: Retrieve all bags that traveled through MIA between 10:00 am and 10:00 pm in February 2019.

```
SELECT bag.bagInfo[].tagNum,bag.bagInfo[].flightLegs[].fltRouteSrc,
$t1 AS HOUR FROM BaggageInfo bag,
EXTRACT(HOUR FROM CAST (bag.bagInfo[0].bagArrivalDate AS Timestamp(0))) $t1,
EXTRACT(YEAR FROM CAST (bag.bagInfo[0].bagArrivalDate AS Timestamp(0))) $t2,
EXTRACT(MONTH FROM CAST (bag.bagInfo[0].bagArrivalDate AS Timestamp(0))) $t3
WHERE bag.bagInfo[].flightLegs[].fltRouteSrc=any "MIA" AND
$t2=2019 AND $t3=02 AND ($t1>10 AND $t1<20)
```

Explanation: You want to know the details of flights that traveled through MIA between 10:00 am and 10:00 pm in February 2019. You use a number of filter conditions here. First, the flight should have originated or traversed through MIA. The year of arrival should be 2019 and the month of arrival should be 2 (February). Then you filter if the hour of arrival is between 10:00 am and 10:00 pm (20 hours).

Output:

```
{"tagNum":"17657806255240","fltRouteSrc":["MIA","LAX"],"HOUR":16}
{"tagNum":"17657806292518","fltRouteSrc":["MIA","LAX"],"HOUR":16}
```

Example 3: Which year and month did the passenger with the reservation code **PQ1M8N** receive the baggage?

```
SELECT fullName,
EXTRACT(YEAR FROM CAST (bag.bagInfo.bagArrivalDate AS Timestamp(0))) AS YEAR,
EXTRACT(MONTH FROM CAST (bag.bagInfo.bagArrivalDate AS Timestamp(0))) AS
MONTH
FROM BaggageInfo bag WHERE bag.confNo="PQ1M8N"
```

Explanation: You first use CAST to convert the `bagArrivalDate` to a TIMESTAMP and then fetch the YEAR component and MONTH component from the Timestamp.

Output:

```
{"fullName":"Kendal Biddle","YEAR":2019,"MONTH":3}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 4: Group the baggage data based on the month of arrival and display the month and the number of baggage that arrived that month.

```
SELECT EXTRACT(MONTH FROM CAST ($bag_arr_date AS Timestamp(0))) AS MONTH,
count(EXTRACT(MONTH FROM CAST ($bag_arr_date AS Timestamp(0)))) AS COUNT
FROM BaggageInfo $bag, $bag.bagInfo[].bagArrivalDate $bag_arr_date
GROUP BY EXTRACT(MONTH FROM CAST ($bag_arr_date AS Timestamp(0)))
```

Explanation: You want to group the data based on the month of the arrival of baggage. You use an unnest array to flatten the data. The `bagInfo` array is flattened and the value of bag arrival date is fetched from the array. You then use CAST to convert the `bagArrivalDate` to a TIMESTAMP and then fetch the YEAR component and MONTH component from the Timestamp. You then use the `count` function to get the total baggage corresponding to every month.

Note

One assumption in the data is that all the baggage has arrived in the same year. So you group the data only based on the month.

Output:

```
{"MONTH":2,"COUNT":11}
{"MONTH":3,"COUNT":10}
```

SQL Expression examples using QueryRequest API

You can use `QueryRequest` API and filter data from a NoSQL table using SQL Expressions.

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***SQLExpressions.java*** from the examples here.

```
 //Fetch rows from the table
private static void fetchRows(NoSQLHandle handle,String sqlstmt) throws
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)){
      for (MapValue res : results) {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                System.out.println("\t" + res);
            }
        }
    }


    String paran_expr="SELECT fullName, bag.bagInfo.tagNum, bag.bagInfo.routing,
    "+
    "bag.bagInfo[].flightLegs[].fltRouteDest FROM BaggageInfo bag WHERE "+
    "bag.bagInfo.flightLegs[].fltRouteSrc=any \"SFO\" AND "+
    "(bag.bagInfo[].flightLegs[].fltRouteDest=any \"ATH\" OR "+
    "bag.bagInfo[].flightLegs[].fltRouteDest=any \"JTR\" )";
    System.out.println("Using Paranthesized expression ");
    fetchRows(handle,paran_expr);

    String case_expr="SELECT fullName,"+
             "CASE WHEN NOT exists bag.bagInfo.flightLegs[0] "+
            "THEN \"you have no bag info\" "+
            "WHEN NOT exists bag.bagInfo.flightLegs[1] "+
            "THEN \"you have one hop\" "+
            "WHEN NOT exists bag.bagInfo.flightLegs[2] "+
            "THEN \"you have two hops.\" "+
            "ELSE \"you have three hops.\" "+
            "END AS NUMBER_HOPS "+
    "FROM BaggageInfo bag WHERE ticketNo=1762341772625";
    System.out.println("Using Case Expression ");
    fetchRows(handle,case_expr);

    String seq_trn_expr="SELECT seq_transform(l.bagInfo[],"+
                                  "seq_transform("+
                                    "$sq1.flightLegs[],"+
                                    "seq_transform("+
                                      "$sq2.actions[],"+
                                      "{"+
                                       "\"at\" : $sq3.actionAt,"+
                                       "\"action\" : $sq3.actionCode,"+
                                       "\"flightNo\" : $sq2.flightNo,"+
                                       "\"tagNum\" : $sq1.tagNum"+
                                      "}"+
                                   ")"+
                                  ")"+
                               ") AS actions FROM baggageInfo l WHERE
    ticketNo=1762376407826";
    System.out.println("Using Sequence Transform Expressions ");
    fetchRows(handle,seq_trn_expr);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***SQLExpressions.py*** from the examples here.

```
# Fetch data from the table
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))


paran_expr = '''SELECT fullName, bag.bagInfo.tagNum, bag.bagInfo.routing,
                   bag.bagInfo[].flightLegs[].fltRouteDest FROM BaggageInfo
bag
                   WHERE bag.bagInfo.flightLegs[].fltRouteSrc=any "SFO" AND
                   (bag.bagInfo[].flightLegs[].fltRouteDest=any "ATH" OR
                   bag.bagInfo[].flightLegs[].fltRouteDest=any "JTR" )'''
print('Using Paranthesized expression:')
fetch_data(handle,paran_expr)

case_expr = '''SELECT fullName,
    CASE
        WHEN NOT exists bag.bagInfo.flightLegs[0]
        THEN "you have no bag info"
        WHEN NOT exists bag.bagInfo.flightLegs[1]
        THEN "you have one hop"
        WHEN NOT exists bag.bagInfo.flightLegs[2]
        THEN "you have two hops."
        ELSE "you have three hops."
    END AS NUMBER_HOPS
    FROM BaggageInfo bag WHERE ticketNo=1762341772625'''
print('Using Case Expression:')
fetch_data(handle,case_expr)

seq_trn_expr = '''SELECT seq_transform(l.bagInfo[],
                     seq_transform(
                       $sq1.flightLegs[],
                         seq_transform(
                           $sq2.actions[],
                            {
                              "at" : $sq3.actionAt,
                              "action" : $sq3.actionCode,
                              "flightNo" : $sq2.flightNo,
                              "tagNum" : $sq1.tagNum
                            }
                        )
                    )
                ) AS actions FROM baggageInfo l WHERE
ticketNo=1762376407826'''
print('Using Sequence Transform Expressions:')
fetch_data(handle,seq_trn_expr)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***SQLExpressions.go*** from the examples here.

```
 //fetch data from the table
func fetchData(client *nosqldb.Client, err error, tableName string, querystmt
string)(){
   prepReq := &nosqldb.PrepareRequest{
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        Statement: querystmt,
    }
    prepRes, err := client.Prepare(prepReq)
    if err != nil {
       fmt.Printf("Prepare failed: %v\n", err)
       return
    }
    queryReq := &nosqldb.QueryRequest{
          PreparedStatement: &prepRes.PreparedStatement,   }
    var results []*types.MapValue
    for {
       queryRes, err := client.Query(queryReq)
       if err != nil {
          fmt.Printf("Query failed: %v\n", err)
       return
       }
       res, err := queryRes.GetResults()
       if err != nil {
          fmt.Printf("GetResults() failed: %v\n", err)
       return
       }
       results = append(results, res...)
       if queryReq.IsDone() {
          break
       }
    }
    for i, r := range results {
       fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
    }
}


paran_expr := `SELECT fullName, bag.bagInfo.tagNum, bag.bagInfo.routing,
                    bag.bagInfo[].flightLegs[].fltRouteDest FROM BaggageInfo
bag
                    WHERE bag.bagInfo.flightLegs[].fltRouteSrc=any "SFO" AND
                    (bag.bagInfo[].flightLegs[].fltRouteDest=any "ATH" OR
                    bag.bagInfo[].flightLegs[].fltRouteDest=any "JTR" )`
fmt.Printf("Using Paranthesized expression:\n")
fetchData(client, err,tableName,paran_expr)

case_expr := `SELECT fullName,
                 CASE
                     WHEN NOT exists bag.bagInfo.flightLegs[0]
                     THEN "you have no bag info"
                     WHEN NOT exists bag.bagInfo.flightLegs[1]
                     THEN "you have one hop"
                     WHEN NOT exists bag.bagInfo.flightLegs[2]
                     THEN "you have two hops."
                     ELSE "you have three hops."
                     END AS NUMBER_HOPS
                     FROM BaggageInfo bag WHERE ticketNo=1762341772625`
fmt.Printf("Using Case Expression:\n")
fetchData(client, err,tableName,case_expr)

seq_trn_expr := `SELECT seq_transform(l.bagInfo[],
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                      seq_transform(
                        $sq1.flightLegs[],
                          seq_transform(
                            $sq2.actions[],
                             {
                               "at" : $sq3.actionAt,
                               "action" : $sq3.actionCode,
                               "flightNo" : $sq2.flightNo,
                               "tagNum" : $sq1.tagNum
                             }
                         )
                     )
                 ) AS actions FROM baggageInfo l WHERE ticketNo=1762376407826`
fmt.Printf("Using Sequence Transform Expressions:\n")
fetchData(client, err,tableName,seq_trn_expr)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***SQLExpressions.js*** from the examples here.

```
  //fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

TypeScript: Download the full code ***SQLExpressions.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: any) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                console.log('  %O', row);
            }
            opt.continuationKey = result.continuationKey;
        } while(opt.continuationKey);
    } catch(error) {
        console.error('  Error: ' + error.message);
    }
}


const paran_expr = `SELECT fullName, bag.bagInfo.tagNum, bag.bagInfo.routing,
                    bag.bagInfo[].flightLegs[].fltRouteDest FROM BaggageInfo
bag
                    WHERE bag.bagInfo.flightLegs[].fltRouteSrc=any "SFO" AND
                    (bag.bagInfo[].flightLegs[].fltRouteDest=any "ATH" OR
                    bag.bagInfo[].flightLegs[].fltRouteDest=any "JTR" )`
console.log("Using Paranthesized expression");
await fetchData(handle,paran_expr);

const case_expr = `SELECT fullName,
                  CASE
                    WHEN NOT exists bag.bagInfo.flightLegs[0]
                    THEN "you have no bag info"
                    WHEN NOT exists bag.bagInfo.flightLegs[1]
                    THEN "you have one hop"
                    WHEN NOT exists bag.bagInfo.flightLegs[2]
                    THEN "you have two hops."
                    ELSE "you have three hops."
                  END AS NUMBER_HOPS
                  FROM BaggageInfo bag WHERE ticketNo=1762341772625`
console.log("Using Case Expression");
await fetchData(handle,case_expr);

const seq_trn_expr = `SELECT seq_transform(l.bagInfo[],
                      seq_transform(
                        $sq1.flightLegs[],
                        seq_transform(
                          $sq2.actions[],
                          {
                            "at" : $sq3.actionAt,
                            "action" : $sq3.actionCode,
                            "flightNo" : $sq2.flightNo,
                            "tagNum" : $sq1.tagNum
                          }
                      )
                  )
                ) AS actions FROM baggageInfo l WHERE ticketNo=1762376407826`
console.log("Using Sequence Transform Expressions");
await fetchData(handle,seq_trn_expr);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***SQLExpressions.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}


private const string paran_expr =@"SELECT fullName, bag.bagInfo.tagNum,
bag.bagInfo.routing,

bag.bagInfo[].flightLegs[].fltRouteDest FROM BaggageInfo bag
                                         WHERE
bag.bagInfo.flightLegs[].fltRouteSrc=any ""SFO"" AND

(bag.bagInfo[].flightLegs[].fltRouteDest=any ""ATH"" OR

bag.bagInfo[].flightLegs[].fltRouteDest=any ""JTR"" )";
Console.WriteLine("\nUsing Paranthesized expression:!");
await fetchData(client,paran_expr);

private const string case_expr =@"SELECT fullName,
                                        CASE
                                          WHEN NOT exists
bag.bagInfo.flightLegs[0]
                                          THEN ""you have no bag info""
                                          WHEN NOT exists
bag.bagInfo.flightLegs[1]
                                          THEN ""you have one hop""
                                          WHEN NOT exists
bag.bagInfo.flightLegs[2]
                                          THEN ""you have two hops.""
                                          ELSE ""you have three hops.""
                                       END AS NUMBER_HOPS
                                       FROM BaggageInfo bag WHERE
ticketNo=1762341772625";
Console.WriteLine("\nUsing Case Expression!");
await fetchData(client,case_expr);

private const string seq_trn_expr =@"SELECT seq_transform(l.bagInfo[],
                                          seq_transform(
                                            $sq1.flightLegs[],
                                            seq_transform(
                                               $sq2.actions[],
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                                                    {
                                                       ""at"" : $sq3.actionAt,

""action"" : $sq3.actionCode,

""flightNo"" : $sq2.flightNo,
                                                       ""tagNum"" : $sq1.tagNum
                                                    }
                                                   )
                                                 )
                                               ) AS actions FROM baggageInfo l
WHERE ticketNo=1762376407826" ;
Console.WriteLine("\nUsing Sequence Transform Expressions!");
await fetchData(client,seq_trn_expr);
```

Timestamp Functions

You can perform various operations on the timestamp and duration values.

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.

The following timestamp functions are supported:

Table 5-1 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.
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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Table 5-1 (Cont.) Timestamp functions

Function	Description
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.

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- Timestamp Arithmetic Functions
- Timestamp Round Functions
- Timestamp Format Functions
- Timestamp Extract Functions
- Current Time Functions
- Examples using QueryRequest API

Timestamp Arithmetic Functions

You can use `timestamp_add`, `timestamp_diff`, or `get_duration` functions to perform arithmetic operations on the timestamp and duration values.

Example 1: 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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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 `flightsLeg` 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 data type.
>
> ISO8601 describes an internationally accepted way to represent dates, times, and durations.
>
> **Syntax:** Date with time: YYYY-MM-DDThh:mm:ss[.s[s[s[s[s[s]]]]][Z|(+|-)hh:mm]
>
> where
>
> - YYYY specifies the year, as four decimal digits
> - MM specifies the month, as two decimal digits, 00 to 12
> - DD specifies the day, as two decimal digits, 00 to 31
> - hh specifies the hour, as two decimal digits, 00 to 23
> - mm specifies the minutes, as two decimal digits, 00 to 59
> - ss[.s[s[s[s[s]]]]] specifies the seconds, as two decimal digits, 00 to 59, optionally followed by a decimal point and 1 to 6 decimal digits (representing the fractional part of a second).
> - Z specifies UTC time (time zone 0). (It can also be specified by +00:00, but not by –00:00.)
> - (+|-)hh:mm specifies the time-zone as difference from UTC. (One of + or – is required.)

Example 1a: 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 on 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 2 : How many 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 week. The `count` function determines the total number of bags in this time period.

Output:

```
{"COUNT_LASTWEEK":0}
```

Example 3: 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.

Output:

```
{"COUNT_NEXT6HOURS":0}
```

Example 4: 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**?

```
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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Explanation: In an airline application every customer can have a 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. To determine the duration in days or hours or minutes, pass the result of the `timestamp_diff` function to the `get_duration` function.

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"}
```

To determine the duration in milliseconds, use the `timestamp_diff` function.

```
SELECT $s.ticketno, $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=1762355527825
```

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 5: 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**?

```
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 action 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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`. To determine the duration in days or hours or minutes, pass the result of the `timestamp_diff` function to the `get_duration` function.

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)
```

Output:

```
{"flightNo":"BM572","checkinTime":"2019-03-02T03:28:00Z",
"bagScanTime":"2019-03-02T04:52:00Z","diff":"- 1 hour 24 minutes"}
```

Example 6: How long does it take for the bags 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 a 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. To determine the duration in days or hours or minutes, pass the result of the `timestamp_diff` function to the `get_duration` function.

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,
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
$bagInfo.flightLegs[0].flightDate) AS diff
FROM baggageinfo $s, $s.bagInfo[] AS $bagInfo
WHERE ticketNo=1762320369957
```

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 Round Functions

You can use `timestamp_ceil`, `timestamp_floor`, `timestamp_trunc`, `timestamp_round`, and `timestamp_bucket` functions to round the timestamp values.

For `timestamp_ceil`, `timestamp_floor`, `timestamp_trunc`, and `timestamp_round` functions, you must supply a `unit` as the second argument. The `unit` specifies the precision to be considered while rounding the input timestamp.

The following units are supported in either singular or plural format: `YEAR, IYEAR, QUARTER, MONTH, WEEK, IWEEK, DAY, HOUR, MINUTE, SECOND`.

You can use the `timestamp_bucket` function to round the given timestamp value to the beginning of the specified interval (bucket). The interval starts at a specified origin on the timeline.

The `timestamp_bucket` supports the following intervals in either singular or plural format: `WEEK, DAY, HOUR, MINUTE, SECOND`.

Example 1: 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"}
```

Example 2: 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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
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 3: 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"}
```

Example 4: 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'
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
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 `timsestamp_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}
```

Timestamp Format Functions

You can use `format_timestamp` and `parse_to_timestamp` functions to format timestamp values. Also, you can use the `to_last_day_of_month` function to fetch the last day of the month from a given timestamp.

Example 1: 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 O", "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"}
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example 2: 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 OOOO"),"yyyy-MM-dd HH:mm:ss OOOO","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"}
```

Example 3: 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

Timestamp extract functions fetch the corresponding date, week, or the index value from a given timestamp.

Date extract functions return the corresponding year/month/day/hour/minute/second/millisecond/microsecond/nanosecond from a timestamp.

Example: 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:

Copyright © 2022, 2025, Oracle and/or its affiliates.

You can use the time functions to retrieve the travel date, month, and year. The `concat` string function is used to concatenate the retrieved travel records to display them in the desired format on the application. You first use the 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 extract functions return the corresponding week/isoweek from a timestamp.

Example: 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 the 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 return the corresponding quarter/week/month/year index from a timestamp.

Example: 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,
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
  "DAYVAL2" : 3
}
```

Current Time Functions

You can use `current_time_millis` and `current_time` functions to fetch the current time. The `current_time_millis` function returns the time as the number of milliseconds. The `current_time` function returns the time as a timestamp value.

Example: 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.

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 the 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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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**.

Examples using QueryRequest API

You can use `QueryRequest` API and apply SQL functions to fetch data from a NoSQL table.

- Java
- Python
- Go
- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***SQLFunctions.java*** from the examples here.

```
 //Fetch rows from the table
private static void fetchRows(NoSQLHandle handle,String sqlstmt) throws 
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)){
      for (MapValue res : results) {
         System.out.println("\t" + res);
      }
   }
}


String ts_func1="SELECT 
timestamp_add(bag.bagInfo.flightLegs[0].estimatedArrival, \"5 minutes\")"+
                         " AS ARRIVAL_TIME FROM BaggageInfo bag WHERE 
ticketNo=1762341772625";
System.out.println("Using timestamp_add function ");
fetchRows(handle,ts_func1);
String 
ts_func2="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=1762344493810";
System.out.println("Using get_duration and timestamp_diff function ");
fetchRows(handle,ts_func2);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***SQLFunctions.py*** from the examples here.

```
# Fetch data from the table
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))


 ts_func1 = '''SELECT
timestamp_add(bag.bagInfo.flightLegs[0].estimatedArrival, "5 minutes")
                 AS ARRIVAL_TIME FROM BaggageInfo bag WHERE
ticketNo=1762341772625'''
print('Using timestamp_add function:')
fetch_data(handle,ts_func1)

ts_func2 =
'''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=1762344493810'''
print('Using get_duration and timestamp_diff function:')
fetch_data(handle,ts_func2)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***SQLFunctions.go*** from the examples here.

```
 //fetch data from the table
func fetchData(client *nosqldb.Client, err error, tableName string, querystmt
string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,   }
   var results []*types.MapValue
   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        return
        }
        res, err := queryRes.GetResults()
        if err != nil {
           fmt.Printf("GetResults() failed: %v\n", err)
        return
        }
        results = append(results, res...)
        if queryReq.IsDone() {
           break
        }
     }
     for i, r := range results {
        fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
     }
  }


ts_func1 := `SELECT timestamp_add(bag.bagInfo.flightLegs[0].estimatedArrival,
"5 minutes")
                AS ARRIVAL_TIME FROM BaggageInfo bag WHERE
ticketNo=1762341772625`
fmt.Printf("Using timestamp_add function::\n")
fetchData(client, err,tableName,ts_func1)

ts_func2 :=
`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=1762344493810`
fmt.Printf("Using get_duration and timestamp_diff function:\n")
fetchData(client, err,tableName,ts_func2)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***SQLFunctions.js*** from the examples here.

```
  //fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    }
}
```

TypeScript: Download the full code ***SQLFunctions.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: string) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}


const ts_func1 = `SELECT
timestamp_add(bag.bagInfo.flightLegs[0].estimatedArrival, "5 minutes")
                  AS ARRIVAL_TIME FROM BaggageInfo bag WHERE
ticketNo=1762341772625`
console.log("Using timestamp_add function:");
await fetchData(handle,ts_func1);

const ts_func2 =
`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=1762344493810`
console.log("Using get_duration and timestamp_diff function:");
await fetchData(handle,ts_func2);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the full code ***SQLFunctions.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}


private const string ts_func1 =@"SELECT
timestamp_add(bag.bagInfo.flightLegs[0].estimatedArrival, ""5 minutes"")
                                       AS ARRIVAL_TIME FROM BaggageInfo bag
WHERE ticketNo=1762341772625";
Console.WriteLine("\nUsing timestamp_add function!");
await fetchData(client,ts_func1);

private const string ts_func2
=@"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=1762344493810";
Console.WriteLine("\nUsing get_duration and timestamp_diff function!");
await fetchData(client,ts_func2);
```

Functions on Strings

There are various built-in functions on strings. In any string, position starts at 0 and ends at length - 1.

If you want to follow along with the examples, see Sample data to run queries to view a sample data and use the scripts to load sample data for testing. The scripts create the tables used in the examples and load data into the tables.

- substring function
- concat function
- upper and lower functions

Copyright © 2022, 2025, Oracle and/or its affiliates.

- trim function
- length function
- contains function
- starts_with and ends_with functions
- index_of function
- replace function
- reverse function
- Examples using QueryRequest API

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.

```
returnvalue substring (source, position [, substring_length] )

source ::= any*
position ::= integer*
substring_length ::= integer*
returnvalue ::= string
```

Example: Fetch the first three characters from the routing details of a passenger with ticket number **1762376407826**.

```
SELECT substring(bag.baginfo.routing,0,3) AS Source
FROM baggageInfo bag
WHERE ticketNo=1762376407826
```

Output:

```
{"Source":"JFK"}
```

concat function

The `concat` function concatenates all its arguments and displays the concatenated string as output.

```
returnvalue concat (source,[source*])
source ::= any*
returnvalue ::= string
```

Example: Display the routing of a customer with a particular ticket number as "The route for passenger_name is ...".

```
SELECT concat("The route for passenger ",fullName , " is ",
bag.baginfo[0].routing)
FROM baggageInfo bag
WHERE ticketNo=1762376407826
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Output:

```
{"Column_1":"The route for passenger Dierdre Amador is JFK/MAD"}
```

upper and lower functions

The `upper` and `lower` are simple functions to convert to fully upper case or lower case respectively. The `upper` function converts all the characters in a string to uppercase. The`lower` function converts all the characters in a string to lowercase.

```
returnvalue upper (source)
returnvalue lower (source)

source ::= any*
returnvalue ::= string
```

Example 1: Fetch the full name of the passenger in uppercase whose ticket number is **1762376407826**.

```
SELECT upper(fullname) AS FULLNAME_CAPITALS
FROM BaggageInfo
WHERE ticketNo=1762376407826
```

Output:

```
{"FULLNAME_CAPITALS":"DIERDRE AMADOR"}
```

Example 2: Fetch the full name of the passenger in lowercase whose ticket number is **1762376407826**.

```
SELECT lower(fullname) AS fullname_lowercase
FROM BaggageInfo WHERE ticketNo=1762376407826
```

Output:

```
{"fullname_lowercase":"dierdre amador"}
```

trim function

The `trim` function enables you to trim leading or trailing characters (or both) from a string. The `ltrim` function enables you to trim leading characters from a string. The `rtrim` function enables you to trim trailing characters from a string.

```
returnvalue trim(source [, position [, trim_character]])

source ::= any*
position ::= "leading"|"trailing"|"both"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
trim_character ::= string*
returnvalue ::= string


returnvalue ltrim(source)

returnvalue rtrim(source)
source ::= any*
returnvalue ::= string
```

Example: Remove leading and trailing blank spaces from the route details of the passenger whose ticket number is **1762350390409**.

```
SELECT trim(bag.baginfo[0].routing,"trailing"," ")
FROM BaggageInfo bag
WHERE ticketNo=1762376407826
```

Output:

```
{"Column_1":"JFK/MAD"}
```

Using `ltrim` function to remove leading spaces:

```
SELECT ltrim(bag.baginfo[0].routing)
FROM BaggageInfo bag
WHERE ticketNo=1762376407826
```

Output:

```
{"Column_1":"JFK/MAD"}
```

Using `rtrim` function to remove trailing spaces:

```
SELECT rtrim(bag.baginfo[0].routing)
FROM BaggageInfo bag
WHERE ticketNo=1762376407826
```

Output:

```
{"Column_1":"JFK/MAD"}
```

length function

The `length` function returns the length of a character string. The length function calculates the length using the UTF character set.

```
returnvalue length(source)

source ::= any*
returnvalue ::= integer
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example: Find the length of the full name of the passenger whose ticket number is **1762350390409**.

```
SELECT fullname, length(fullname) AS fullname_length
FROM BaggageInfo
WHERE ticketNo=1762350390409
```

Output:

```
{"fullname":"Fallon Clements","fullname_length":15}
```

contains function

The `contains` function indicates whether or not a search string is present inside the source string.

```
returnvalue contains(source, search_string)

source ::= any*
search_string ::= any*
returnvalue ::= boolean
```

Example: Fetch the full names of passengers who have "SFO" in their route.

```
SELECT fullname FROM baggageInfo bag
WHERE EXISTS bag.bagInfo[contains($element.routing,"SFO")]
```

Output:

```
{"fullname":"Michelle Payne"}
{"fullname":"Lucinda Beckman"}
{"fullname":"Henry Jenkins"}
{"fullname":"Lorenzo Phil"}
{"fullname":"Gerard Greene"}
```

starts_with and ends_with functions

The `starts_with` function indicates whether or not the source string begins with the search string.

```
returnvalue starts_with(source, search_string)

source ::= any*
search_string ::= any*
returnvalue ::= boolean
```

The `ends_with`function indicates whether or not the source string ends with the search string.

```
returnvalue ends_with(source, search_string)

source ::= any*
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
search_string ::= any*
returnvalue ::= boolean
```

Example: 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**?

```
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 2 : Find list of passengers whose destination is **JTR**.

```
SELECT fullname FROM baggageInfo $bagInfo
WHERE ends_with($bagInfo.bagInfo[].routing, "JTR")
```

Output:

```
{"fullname":"Lucinda Beckman"}
{"fullname":"Gerard Greene"}
{"fullname":"Michelle Payne"}
```

index_of function

The `index_of` function determines the position of the first character of the search string at its first occurrence if any.

```
returnvalue index_of(source, search_string [, start_position])
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
source ::= any*
search_string ::= any*
start_position ::= integer*
returnvalue ::= integer
```

Various return values:

- Returns the position of the first character of the search string at its first occurrence. The position is relative to the start position of the string (which is zero).
- Returns -1 if `search_string` is not present in the 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.

Example 1: Determine at which position "-" is found in the estimated arrival time of the first leg for the passenger with ticket number **1762320569757**.

```
SELECT index_of(bag.baginfo.flightLegs[0].estimatedArrival,"-")
FROM BaggageInfo bag
WHERE ticketNo=1762320569757
```

Output:

```
{"Column_1":4}
```

Example 2: Determine at which position "/" is found in the routing of the first leg for passenger with ticket number **1762320569757**. This will help you determine how many characters are there for the source point for the passenger with ticket number **1762320569757**.

```
SELECT index_of(bag.baginfo.routing,"/")
FROM BaggageInfo bag
WHERE ticketNo=1762320569757
```

Output:

```
"Column_1":3}
```

replace function

The `replace` function returns the source with every occurrence of the search string replaced with the replacement string.

```
returnvalue replace(source, search_string [, replacement_string])

source ::= any*
search_string ::= any*
replacement_string ::= any*
returnvalue ::= string
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Example: Replace the source location of the passenger with ticket number **1762320569757** from SFO to **SOF**.

```
SELECT replace(bag.bagInfo[0].routing,"SFO","SOF")
FROM baggageInfo bag
WHERE ticketNo=1762320569757
```

Output:

```
{"Column_1":"SOF/IST/ATH/JTR"}
```

Example 2: Replace the double quote in the passenger name with a single quote.

If your data might contain a double quote in the passenger's name, you can use replace function to change the double quote to a single quote.

```
SELECT fullname,
replace(fullname, "\"", "'") as new_fullname
FROM BaggageInfo bag
```

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.

```
returnvalue reverse(source)

source ::= any*
returnvalue ::= string
```

Example: Display the full name and reverse the full name of the passenger with ticket number **1762330498104**.

```
SELECT fullname, reverse(fullname)
FROM baggageInfo
WHERE ticketNo=1762330498104
```

Output:

```
{"fullname":"Michelle Payne","Column_2":"enyaP ellehciM"}
```

Examples using QueryRequest API

You can use `QueryRequest` API and apply SQL functions to fetch data from a NoSQL table.

- Java
- Python
- Go

Copyright © 2022, 2025, Oracle and/or its affiliates.

- Node.js
- C#

Java

To execute your query, you use the `NoSQLHandle.query()` API.

Download the full code ***SQLFunctions.java*** from the examples here.

```
 //Fetch rows from the table
private static void fetchRows(NoSQLHandle handle,String sqlstmt) throws 
Exception {
   try (
      QueryRequest queryRequest = new QueryRequest().setStatement(sqlstmt);
      QueryIterableResult results = handle.queryIterable(queryRequest)){
      for (MapValue res : results) {
         System.out.println("\t" + res);
      }
   }
}


String string_func1="SELECT substring(bag.baginfo.routing,0,3) AS Source FROM 
baggageInfo bag WHERE ticketNo=1762376407826";
System.out.println("Using substring function ");
fetchRows(handle,string_func1);
String string_func2="SELECT fullname, length(fullname) AS fullname_length 
FROM BaggageInfo WHERE ticketNo=1762320369957";
System.out.println("Using length function ");
fetchRows(handle,string_func2);
String string_func3="SELECT fullname FROM baggageInfo bag WHERE EXISTS 
bag.bagInfo[contains($element.routing,\"SFO\")]";
System.out.println("Using contains function ");
fetchRows(handle,string_func3);
```

Python

To execute your query use the `borneo.NoSQLHandle.query()` method.

Download the full code ***SQLFunctions.py*** from the examples here.

```
# Fetch data from the table
def fetch_data(handle,sqlstmt):
   request = QueryRequest().set_statement(sqlstmt)
   print('Query results for: ' + sqlstmt)
   result = handle.query(request)
   for r in result.get_results():
      print('\t' + str(r))



 string_func1 = '''SELECT substring(bag.baginfo.routing,0,3) AS Source FROM 
baggageInfo bag 
                     WHERE ticketNo=1762376407826'''
 print('Using substring function:')
 fetch_data(handle,string_func1)
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
 string_func2 = '''SELECT fullname, length(fullname) AS fullname_length FROM 
BaggageInfo
                     WHERE ticketNo=1762320369957'''
 print('Using length function:')
 fetch_data(handle,string_func2)

 string_func3 = '''SELECT fullname FROM baggageInfo bag WHERE
                     EXISTS bag.bagInfo[contains($element.routing,"SFO")]'''
 print('Using contains function:')
 fetch_data(handle,string_func3)
```

Go

To execute a query use the `Client.Query` function.

Download the full code ***SQLFunctions.go*** from the examples here.

```
 //fetch data from the table
func fetchData(client *nosqldb.Client, err error, tableName string, querystmt 
string)(){
   prepReq := &nosqldb.PrepareRequest{
        Statement: querystmt,
   }
   prepRes, err := client.Prepare(prepReq)
   if err != nil {
      fmt.Printf("Prepare failed: %v\n", err)
      return
   }
   queryReq := &nosqldb.QueryRequest{
         PreparedStatement: &prepRes.PreparedStatement,    }
   var results []*types.MapValue
   for {
      queryRes, err := client.Query(queryReq)
      if err != nil {
         fmt.Printf("Query failed: %v\n", err)
      return
      }
      res, err := queryRes.GetResults()
      if err != nil {
         fmt.Printf("GetResults() failed: %v\n", err)
      return
      }
      results = append(results, res...)
      if queryReq.IsDone() {
         break
      }
   }
   for i, r := range results {
      fmt.Printf("\t%d: %s\n", i+1, jsonutil.AsJSON(r.Map()))
   }
}


string_func1 := `SELECT substring(bag.baginfo.routing,0,3) AS Source FROM 
baggageInfo bag
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                    WHERE ticketNo=1762376407826`
fmt.Printf("Using substring function:\n")
fetchData(client, err,tableName,string_func1)

string_func2 := `SELECT fullname, length(fullname) AS fullname_length FROM
BaggageInfo
                    WHERE ticketNo=1762320369957`
fmt.Printf("Using length function:\n")
fetchData(client, err,tableName,string_func2)

string_func3 := `SELECT fullname FROM baggageInfo bag WHERE
                    EXISTS bag.bagInfo[contains($element.routing,"SFO")]`
fmt.Printf("Using contains function:\n")
fetchData(client, err,tableName,string_func3)
```

Node.js

To execute a query use `query` method.

JavaScript: Download the full code ***SQLFunctions.js*** from the examples here.

```
  //fetches data from the table
async function fetchData(handle,querystmt) {
   const opt = {};
   try {
      do {
         const result = await handle.query(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
         }
         opt.continuationKey = result.continuationKey;
      } while(opt.continuationKey);
   } catch(error) {
      console.error('  Error: ' + error.message);
   }
}
```

TypeScript: Download the full code ***SQLFunctions.ts*** from the examples here.

```
interface StreamInt {
   acct_Id: Integer;
   profile_name: String;
   account_expiry: TIMESTAMP;
   acct_data: JSON;
}


/* fetches data from the table */
async function fetchData(handle: NoSQLClient,querystmt: any) {
   const opt = {};
   try {
      do {
         const result = await handle.query<StreamInt>(querystmt, opt);
         for(let row of result.rows) {
            console.log('  %O', row);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            }
            opt.continuationKey = result.continuationKey;
         } while(opt.continuationKey);
      } catch(error) {
         console.error('  Error: ' + error.message);
      }
   }


const string_func1 = `SELECT substring(bag.baginfo.routing,0,3) AS Source
FROM baggageInfo bag
                      WHERE ticketNo=1762376407826`
console.log("Using substring function:");
await fetchData(handle,string_func1);

const string_func2 = `SELECT fullname, length(fullname) AS fullname_length
FROM BaggageInfo
                      WHERE ticketNo=1762320369957`
console.log("Using length function");
await fetchData(handle,string_func2);

const string_func3 = `SELECT fullname FROM baggageInfo bag WHERE
                      EXISTS bag.bagInfo[contains($element.routing,"SFO")]`
console.log("Using contains function");
await fetchData(handle,string_func3);
```

C#

To execute a query, you may call `QueryAsync` method or call `GetQueryAsyncEnumerable` method and iterate over the resulting async enumerable.

Download the full code ***SQLFunctions.cs*** from the examples here.

```
private static async Task fetchData(NoSQLClient client,String querystmt){
   var queryEnumerable = client.GetQueryAsyncEnumerable(querystmt);
   await DoQuery(queryEnumerable);
}

private static async Task DoQuery(IAsyncEnumerable<QueryResult<RecordValue>>
queryEnumerable){
   Console.WriteLine("  Query results:");
   await foreach (var result in queryEnumerable) {
      foreach (var row in result.Rows)
      {
         Console.WriteLine();
         Console.WriteLine(row.ToJsonString());
      }
   }
}


private const string string_func1 =@"SELECT
substring(bag.baginfo.routing,0,3) AS Source FROM baggageInfo bag
                                          WHERE ticketNo=1762376407826" ;
Console.WriteLine("\nUsing substring function!");
await fetchData(client,string_func1);
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
private const string string_func2 =@"SELECT fullname, length(fullname) AS
fullname_length FROM BaggageInfo
                                            WHERE ticketNo=1762320369957";
Console.WriteLine("\nUsing length function!");
await fetchData(client,string_func2);


private const string string_func3 =@"SELECT fullname FROM baggageInfo bag
WHERE
                                            EXISTS
bag.bagInfo[contains($element.routing,""SFO"")]";
Console.WriteLine("\nUsing contains function!");
await fetchData(client,string_func3);
```

Query execution plan

A query execution plan is the sequence of operations Oracle NoSQL Database performs to run a query.

- Overview of query plan
- Query 1: Using primary key index with an index range scan
- Query 2: Using primary key index with an index predicate
- Query 3: Using a secondary index with an index range scan
- Query 4: Using the primary index
- Query 5: Sort the data using a Covering index
- Query 6: Using a secondary index with an index predicate
- Query 7: Group data with fields as part of the index
- Query 8: Using the secondary index with multiple index scans
- Query 9: A SINGLE PARTITION query using a primary index
- Query 10: Group data with fields not part of any index

Overview of query plan

A query execution plan is internally 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 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.

Some of the most common and important iterators used in queries are :

Copyright © 2022, 2025, Oracle and/or its affiliates.

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
- Retrieve 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:

1. The RECEIVE iterator itself and all iterators that are above it in the iterator tree are executed at the driver.
2. 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.
- **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.

Populating the tables to view the query execution plan :

If you want to follow along with the examples, download the script baggageschema_loaddata.sql and run it as shown below. This script creates the table used in the example 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 baggageschema_loaddata.sql
```

Creating indexes:

Create the following indexes in the `baggageInfo` table as shown below.

1. Create an index on passengers reservation code.

   ```
   CREATE INDEX fixedschema_conf ON baggageInfo confNo)
   ```

2. Create an index on the full name and phone number of passengers

   ```
   CREATE INDEX compindex_namephone ON baggageInfo (fullName,contactPhone)
   ```

3. Create an index on three fields, when the bag was last seen, the last seen station, and the arrival date and time.

   ```
   CREATE INDEX simpleindex_arrival ON
   baggageInfo(bagInfo[].lastSeenTimeGmt as ANYATOMIC,
   ```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    bagInfo[].bagArrivalDate as ANYATOMIC,
    bagInfo[].lastSeenTimeStation as ANYATOMIC)
```

Query 1: Using primary key index with an index range scan

Fetch the bag details of passengers for ticket numbers in a range.

```
SELECT fullname, ticketNo,
bag.bagInfo[].tagNum,bag.bagInfo[].routing
FROM BaggageInfo bag WHERE
1762340000000 < ticketNo AND ticketNo < 1762352000000
```

Plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_PARTITIONS",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "BaggageInfo",
      "row variable" : "$$bag",
      "index used" : "primary index",
      "covering index" : false,
      "index scans" : [
        {
          "equality conditions" : {},
          "range conditions" : { "ticketNo" : { "start value" : 
1762340000000, "start inclusive" : false, "end value" : 1762352000000, "end 
inclusive" : false } }
        }
      ]
    },
    "FROM variable" : "$$bag",
    "SELECT expressions" : [
      {
        "field name" : "fullname",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "fullname",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
          }
        }
      },
      {
        "field name" : "ticketNo",
        "field expression" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "ticketNo",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$bag"
      }
    }
  },
  {
    "field name" : "tagNum",
    "field expression" :
    {
      "iterator kind" : "ARRAY_CONSTRUCTOR",
      "conditional" : true,
      "input iterators" : [
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "tagNum",
          "input iterator" :
          {
            "iterator kind" : "ARRAY_FILTER",
            "input iterator" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "bagInfo",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$$bag"
              }
            }
          }
        }
      ]
    }
  },
  {
    "field name" : "routing",
    "field expression" :
    {
      "iterator kind" : "ARRAY_CONSTRUCTOR",
      "conditional" : true,
      "input iterators" : [
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "routing",
          "input iterator" :
          {
            "iterator kind" : "ARRAY_FILTER",
            "input iterator" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "bagInfo",
              "input iterator" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                  {
                    "iterator kind" : "VAR_REF",
                    "variable" : "$$bag"
                  }
                }
              }
            }
          ]
        }
      }
    ]
  }
}
```

Explanation:

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator.
- The value of the `FROM` field is an iterator; in this case, it is a **TABLE** iterator.
- The **primary key index** is used and the index is not covering (as you need to scan the table to fetch columns other than the index entries).
- The index scan property contains the start and stop conditions that define the index scans to be performed.
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable (`$$bag`)is the same as the **row variable** of the TABLE iterator, as the index used is not covering.
- In the SELECT expression four fields (`fullname,ticketNo,bag.bagInfo[].tagNum,bag.bagInfo[].routing`) are fetched. These correspond to four field names and field expressions in the SELECT expression clause. For the first two fields, the field expression is computed using `FIELD_STEP` iterator. For the last 2 fields, an `ARRAY_CONSTRUCTOR` iterator is used which iterates over the corresponding arrays to fetch the field value.

Query 2: Using primary key index with an index predicate

Fetch the bag details of passengers for ticket numbers satisfying one of the two ranges of values.

```
SELECT fullname, ticketNo, bag.bagInfo[].tagNum,bag.bagInfo[].routing
FROM BaggageInfo bag WHERE ticketNo > 1762340000000 OR ticketNo <
1762352000000;
```

Plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_PARTITIONS",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{
  "iterator kind" : "TABLE",
  "target table" : "BaggageInfo",
  "row variable" : "$$bag",
  "index used" : "primary index",
  "covering index" : false,
  "index scans" : [
    {
      "equality conditions" : {},
      "range conditions" : {}
    }
  ],
  "index filtering predicate" :
  {
    "iterator kind" : "OR",
    "input iterators" : [
      {
        "iterator kind" : "GREATER_THAN",
        "left operand" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "ticketNo",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
          }
        },
        "right operand" :
        {
          "iterator kind" : "CONST",
          "value" : 1762340000000
        }
      },
      {
        "iterator kind" : "LESS_THAN",
        "left operand" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "ticketNo",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
          }
        },
        "right operand" :
        {
          "iterator kind" : "CONST",
          "value" : 1762352000000
        }
      }
    ]
  }
},
"FROM variable" : "$$bag",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
"SELECT expressions" : [
  {
    "field name" : "fullname",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "fullname",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$bag"
      }
    }
  },
  {
    "field name" : "ticketNo",
    "field expression" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "ticketNo",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$bag"
      }
    }
  },
  {
    "field name" : "tagNum",
    "field expression" :
    {
      "iterator kind" : "ARRAY_CONSTRUCTOR",
      "conditional" : true,
      "input iterators" : [
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "tagNum",
          "input iterator" :
          {
            "iterator kind" : "ARRAY_FILTER",
            "input iterator" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "bagInfo",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$$bag"
              }
            }
          }
        }
      ]
    }
  },
  {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            "field name" : "routing",
            "field expression" :
            {
              "iterator kind" : "ARRAY_CONSTRUCTOR",
              "conditional" : true,
              "input iterators" : [
                {
                  "iterator kind" : "FIELD_STEP",
                  "field name" : "routing",
                  "input iterator" :
                  {
                    "iterator kind" : "ARRAY_FILTER",
                    "input iterator" :
                    {
                      "iterator kind" : "FIELD_STEP",
                      "field name" : "bagInfo",
                      "input iterator" :
                      {
                        "iterator kind" : "VAR_REF",
                        "variable" : "$$bag"
                      }
                    }
                  }
                }
              ]
            }
          }
        ]
      }
    }
```

Explanation:

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator.
- The value of the **`FROM`** field is an iterator; in this case, it is a **TABLE** iterator.
- The **primary key index** is used and the index is not covering (as you need to scan the table to fetch columns other than the index entries).
- The **`index filtering predicate`** evaluates the filter criteria on the **`ticketNo`** field. Using the greater than and less than operators the filter condition is evaluated.
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable (**`$$bag`**) is the same as the **row variable** of the TABLE iterator, as the index used is not covering.
- In the SELECT expression four fields (**`fullname,ticketNo,bag.bagInfo[].tagNum, bag.bagInfo[].routing`**) are fetched. These correspond to four field names and field expressions in the SELECT expression clause. For the first two fields, the field expression is computed using **`FIELD_STEP`** iterator. For the last 2 fields, an **`ARRAY_CONSTRUCTOR`** iterator is used which iterates over the corresponding arrays to fetch the field value.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Query 3: Using a secondary index with an index range scan

Fetch the bag details for a particular reservation code.

```
SELECT fullName,bag.ticketNo, bag.confNo, bag.bagInfo[].tagNum,
bag.bagInfo[].routing FROM BaggageInfo bag WHERE bag.confNo="FH7G1W"
```

Plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_SHARDS",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "BaggageInfo",
      "row variable" : "$$bag",
      "index used" : "fixedschema_conf",
      "covering index" : false,
      "index scans" : [
        {
          "equality conditions" : {"confNo":"FH7G1W"},
          "range conditions" : {}
        }
      ]
    },
    "FROM variable" : "$$bag",
    "SELECT expressions" : [
      {
        "field name" : "fullName",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "fullName",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
          }
        }
      },
      {
        "field name" : "ticketNo",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "ticketNo",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
        }
      }
    },
    {
      "field name" : "confNo",
      "field expression" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "confNo",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$bag"
        }
      }
    },
    {
      "field name" : "tagNum",
      "field expression" :
      {
        "iterator kind" : "ARRAY_CONSTRUCTOR",
        "conditional" : true,
        "input iterators" : [
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "tagNum",
            "input iterator" :
            {
              "iterator kind" : "ARRAY_FILTER",
              "input iterator" :
              {
                "iterator kind" : "FIELD_STEP",
                "field name" : "bagInfo",
                "input iterator" :
                {
                  "iterator kind" : "VAR_REF",
                  "variable" : "$$bag"
                }
              }
            }
          }
        ]
      }
    },
    {
      "field name" : "routing",
      "field expression" :
      {
        "iterator kind" : "ARRAY_CONSTRUCTOR",
        "conditional" : true,
        "input iterators" : [
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "routing",
            "input iterator" :
            {
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                    "iterator kind" : "ARRAY_FILTER",
                    "input iterator" :
                    {
                      "iterator kind" : "FIELD_STEP",
                      "field name" : "bagInfo",
                      "input iterator" :
                      {
                        "iterator kind" : "VAR_REF",
                        "variable" : "$$bag"
                      }
                    }
                  }
                }
              ]
            }
          }
        ]
      }
}
```

Explanation:

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator.
- The value of the `FROM` field is an iterator; in this case, it is a **TABLE** iterator.
- The index `fixedschema_conf` is used and the index is not covering (as you need to scan the table to fetch columns other than the index entries).
- The index scan property contains the start and stop conditions that define the index scans to be performed.
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable (`$$bag`) is the same as the **row variable** of the TABLE iterator, as the index used is not covering.
- In the SELECT expression four fields (`fullname,ticketNo,confNo,bag.bagInfo[].tagNum,bag.bagInfo[].routing`) are fetched. These correspond to five field names and field expressions in the SELECT expression clause. For the first three fields, the field expression is computed using `FIELD_STEP` iterator. For the last two fields, an `ARRAY_CONSTRUCTOR` iterator is used which iterates over the corresponding arrays to fetch the field value.

Query 4: Using the primary index

Fetch the name and routing details of all male passengers.

```
SELECT fullname,bag.bagInfo[].routing FROM BaggageInfo bag
WHERE gender!="F"
```

Plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_PARTITIONS",
  "input iterator" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
{
  "iterator kind" : "SELECT",
  "FROM" :
  {
    "iterator kind" : "TABLE",
    "target table" : "BaggageInfo",
    "row variable" : "$$bag",
    "index used" : "primary index",
    "covering index" : false,
    "index scans" : [
      {
        "equality conditions" : {},
        "range conditions" : {}
      }
    ]
  },
  "FROM variable" : "$$bag",
  "WHERE" :
  {
    "iterator kind" : "NOT_EQUAL",
    "left operand" :
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "gender",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$$bag"
      }
    },
    "right operand" :
    {
      "iterator kind" : "CONST",
      "value" : "F"
    }
  },
  "SELECT expressions" : [
    {
      "field name" : "fullname",
      "field expression" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "fullname",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$bag"
        }
      }
    },
    {
      "field name" : "routing",
      "field expression" :
      {
        "iterator kind" : "ARRAY_CONSTRUCTOR",
        "conditional" : true,
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
              "input iterators" : [
                {
                  "iterator kind" : "FIELD_STEP",
                  "field name" : "routing",
                  "input iterator" :
                  {
                    "iterator kind" : "ARRAY_FILTER",
                    "input iterator" :
                    {
                      "iterator kind" : "FIELD_STEP",
                      "field name" : "bagInfo",
                      "input iterator" :
                      {
                        "iterator kind" : "VAR_REF",
                        "variable" : "$$bag"
                      }
                    }
                  }
                }
              ]
            }
          }
        ]
      }
    }
```

Explanation:

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator.
- The value of the `FROM` field is an iterator; in this case, it is a **TABLE** iterator.
- The **primary key index** is used and the index is not covering (as you need to scan the table to fetch columns other than the index entries).
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable (`$$bag`) is the same as the **row variable** of the TABLE iterator, as the index used is not covering.
- In the SELECT expression two fields (`fullname,bag.bagInfo[].routing`) are fetched. These correspond to two field names and field expressions in the SELECT expression clause. For the first field, the field expression is computed using `FIELD_STEP` iterator. For the second field, an `ARRAY_CONSTRUCTOR` iterator is used which iterates over the corresponding array to fetch the field value.

Query 5: Sort the data using a Covering index

Fetch the name and phone number of all passengers.

```
SELECT bag.contactPhone, bag.fullName FROM BaggageInfo bag
ORDER BY bag.fullName
```

Plan:

```
{
  "iterator kind" : "RECEIVE",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
    "distribution kind" : "ALL_SHARDS",
    "order by fields at positions" : [ 1 ],
    "input iterator" :
    {
      "iterator kind" : "SELECT",
      "FROM" :
      {
        "iterator kind" : "TABLE",
        "target table" : "BaggageInfo",
        "row variable" : "$$bag",
        "index used" : "compindex_namephone",
        "covering index" : true,
        "index row variable" : "$$bag_idx",
        "index scans" : [
          {
            "equality conditions" : {},
            "range conditions" : {}
          }
        ]
      },
      "FROM variable" : "$$bag_idx",
      "SELECT expressions" : [
        {
          "field name" : "contactPhone",
          "field expression" :
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "contactPhone",
            "input iterator" :
            {
              "iterator kind" : "VAR_REF",
              "variable" : "$$bag_idx"
            }
          }
        },
        {
          "field name" : "fullName",
          "field expression" :
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "fullName",
            "input iterator" :
            {
              "iterator kind" : "VAR_REF",
              "variable" : "$$bag_idx"
            }
          }
        }
      ]
    }
  }
```

Explanation:

Copyright © 2022, 2025, Oracle and/or its affiliates.

- 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 **`fullName`**. The `fullName` is part of the **`compindex_namephone`** 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 which is sorted in the SELECT expression.

  ```
  "order by fields at positions" : [ 1 ]
  ```

- In this example, the order by is done using the **`fullName`** 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.
- The index **`compindex_namephone`** 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 **`$$bag_idx`** which is the name of a variable ranging over the index entries produced by the TABLE iterator. Every time a new index entry is produced by the index scan, the `$$bag_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 the same as the index row variable (**`$$bag_idx`**) as the index is covering.
- This index row variable (`$$bag_idx`) will be referenced by iterators implementing the other clauses of the SELECT expression.
- In the SELECT expression two fields (`contactPhone,fullName`) are fetched. These correspond to two field names and field expressions in the SELECT expression clause. For both fields, the field expression is computed using **`FIELD_STEP`** iterator.

Query 6: Using a secondary index with an index predicate

Fetch the name, ticket number, and arrival date of passengers whose arrival date is greater than a given value.

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

Plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "ALL_SHARDS",
  "distinct by fields at positions" : [ 1 ],
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "BaggageInfo",
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
      "row variable" : "$$bag",
      "index used" : "simpleindex_arrival",
      "covering index" : false,
      "index row variable" : "$$bag_idx",
      "index scans" : [
        {
          "equality conditions" : {},
          "range conditions" : {}
        }
      ],
      "index filtering predicate" :
      {
        "iterator kind" : "GREATER_OR_EQUAL",
        "left operand" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "bagInfo[].bagArrivalDate",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag_idx"
          }
        },
        "right operand" :
        {
          "iterator kind" : "CONST",
          "value" : "2019-01-01T00:00:00"
        }
      }
    },
    "FROM variable" : "$$bag",
    "SELECT expressions" : [
      {
        "field name" : "fullName",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "fullName",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
          }
        }
      },
      {
        "field name" : "ticketNo",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "ticketNo",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$$bag"
          }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
          }
        },
        {
          "field name" : "bagArrivalDate",
          "field expression" :
          {
            "iterator kind" : "ARRAY_CONSTRUCTOR",
            "conditional" : true,
            "input iterators" : [
              {
                "iterator kind" : "FIELD_STEP",
                "field name" : "bagArrivalDate",
                "input iterator" :
                {
                  "iterator kind" : "ARRAY_FILTER",
                  "input iterator" :
                  {
                    "iterator kind" : "FIELD_STEP",
                    "field name" : "bagInfo",
                    "input iterator" :
                    {
                      "iterator kind" : "VAR_REF",
                      "variable" : "$$bag"
                    }
                  }
                }
              }
            ]
          }
        }
      ]
    }
  }
```

Explanation:

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator.
- The value of the **FROM** field is an iterator; in this case, it is a **TABLE** iterator.
- The EXISTS condition is actually converted to a filtering predicate. There is one filtering predicate which is the whole WHERE expression. The index **simpleindex_arrival** is the only one applicable here and is used.
- The **index filtering predicate** evaluates the filter criteria on the **bagArrivalDate** field. Using the greater than and less than operators the filter condition is evaluated.
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable (**$$bag**)is the same as the **row variable** of the TABLE iterator, as the index used is not covering.
- In the SELECT expression three fields (**fullname, ticketNo,bag.bagInfo[].bagArrivalDat**) are fetched. These correspond to three field names and field expressions in the SELECT expression clause. For the first two fields, the field expression is computed using **FIELD_STEP** iterator. For the last field, an

Copyright © 2022, 2025, Oracle and/or its affiliates.

ARRAY_CONSTRUCTOR iterator is used which iterates over the corresponding arrays to fetch the field value.

Query 7: Group data with fields as part of the index

Fetch the reservation code and count of bags for all passengers.

```
SELECT bag.confNo, count(bag.bagInfo) AS TOTAL_BAGS
FROM BaggageInfo bag GROUP BY bag.confNo;
```

Plan:

```
{
  "iterator kind" : "SELECT",
  "FROM" :
  {
    "iterator kind" : "RECEIVE",
    "distribution kind" : "ALL_SHARDS",
    "order by fields at positions" : [ 0 ],
    "input iterator" :
    {
      "iterator kind" : "SELECT",
      "FROM" :
      {
        "iterator kind" : "TABLE",
        "target table" : "BaggageInfo",
        "row variable" : "$$bag",
        "index used" : "fixedschema_conf",
        "covering index" : false,
        "index scans" : [
          {
            "equality conditions" : {},
            "range conditions" : {}
          }
        ]
      },
      "FROM variable" : "$$bag",
      "GROUP BY" : "Grouping by the first expression in the SELECT list",
      "SELECT expressions" : [
        {
          "field name" : "confNo",
          "field expression" :
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "confNo",
            "input iterator" :
            {
              "iterator kind" : "VAR_REF",
              "variable" : "$$bag"
            }
          }
        },
        {
          "field name" : "TOTAL_BAGS",
          "field expression" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
            {
              "iterator kind" : "FN_COUNT",
              "input iterator" :
              {
                "iterator kind" : "FIELD_STEP",
                "field name" : "bagInfo",
                "input iterator" :
                {
                  "iterator kind" : "VAR_REF",
                  "variable" : "$$bag"
                }
              }
            }
          }
        ]
      }
    },
    "FROM variable" : "$from-1",
    "GROUP BY" : "Grouping by the first expression in the SELECT list",
    "SELECT expressions" : [
      {
        "field name" : "confNo",
        "field expression" :
        {
          "iterator kind" : "FIELD_STEP",
          "field name" : "confNo",
          "input iterator" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$from-1"
          }
        }
      },
      {
        "field name" : "TOTAL_BAGS",
        "field expression" :
        {
          "iterator kind" : "FUNC_SUM",
          "input iterator" :
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "TOTAL_BAGS",
            "input iterator" :
            {
              "iterator kind" : "VAR_REF",
              "variable" : "$from-1"
            }
          }
        }
      }
    ]
  }
```

Explanation:

Copyright © 2022, 2025, Oracle and/or its affiliates.

- In this query, you group all bags based on the **confNo** of the users and determine the total count of bags belonging to each **confNo**.
- The group-by is index-based, that is the group-by field (confNo) is 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 is the first expression in the SELECT list (**bag.confNo**).

  ```
  "GROUP BY" : "Grouping by the first expression in the SELECT list"
  ```

- The index **fixedschema_conf** is used here and in this example, it is a non-covering index as the query also needs to fetch **count(bag.bagInfo)**which is outside of the entries of the index.
- 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 inner FROM variable is the same as the row variable (**$$bag**) as the index is not covering.
- In the SELECT expression two fields are fetched: **bag.confNo,count(bag.bagInfo)**. These correspond to two 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 bag.confNo may exist at multiple RNs. So, regrouping and re-aggregation have to be performed by the driver. This is done by the outer SELECT iterator (above the RECEIVE iterator).
- The result is also sorted by confNo. 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 which is sorted in the SELECT expression. In this example **bag.confNo** is the first field in the SELECT expression. So order by fields at positions takes an array index of 0.

  ```
  "order by fields at positions" : [ 0 ]
  ```

- In the outer SELECT expression, two fields are fetched: bag.confNo,count(bag.bagInfo). The $from-1 FROM variable will be referenced by iterators implementing the other clauses of the outer SELECT expression. These correspond to two field names and field expressions in the outer SELECT expression clause. For the first field, the field expression uses **FIELD_STEP** iterator. The second field is the aggregate function **count**. The iterator **FUNC_SUM**is used to iterate the result produced by its parent iterator and determine the total number of bags.

Query 8: Using the secondary index with multiple index scans

Fetch the full name and tag number of passengers who are in the given list of names.

```
SELECT bagdet.fullName, bagdet.bagInfo[].tagNum
FROM BaggageInfo bagdet WHERE bagdet.fullName IN
("Lucinda Beckman", "Adam Phillips",
"Zina Christenson","Fallon Clements");
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Plan:

```
{
  "iterator kind" : "SELECT",
  "FROM" :
  {
     "iterator kind" : "RECEIVE",
     "distribution kind" : "ALL_SHARDS",
     "order by fields at positions" : [ 0 ],
     "input iterator" :
     {
       "iterator kind" : "SELECT",
       "FROM" :
       {
         "iterator kind" : "TABLE",
         "target table" : "BaggageInfo",
         "row variable" : "$$bag",
         "index used" : "fixedschema_conf",
         "covering index" : false,
         "index scans" : [
           {
             "equality conditions" : {},
             "range conditions" : {}
           }
         ]
      },
      "FROM variable" : "$$bag",
      "GROUP BY" : "Grouping by the first expression in the SELECT list",
      "SELECT expressions" : [
         {
           "field name" : "confNo",
           "field expression" :
           {
             "iterator kind" : "FIELD_STEP",
             "field name" : "confNo",
             "input iterator" :
             {
               "iterator kind" : "VAR_REF",
               "variable" : "$$bag"
             }
           }
         },
         {
           "field name" : "TOTAL_BAGS",
           "field expression" :
           {
             "iterator kind" : "FN_COUNT",
             "input iterator" :
             {
               "iterator kind" : "FIELD_STEP",
               "field name" : "bagInfo",
               "input iterator" :
               {
                 "iterator kind" : "VAR_REF",
                 "variable" : "$$bag"
               }
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                }
              }
            }
          ]
        }
      },
      "FROM variable" : "$from-1",
      "GROUP BY" : "Grouping by the first expression in the SELECT list",
      "SELECT expressions" : [
        {
          "field name" : "confNo",
          "field expression" :
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "confNo",
            "input iterator" :
            {
              "iterator kind" : "VAR_REF",
              "variable" : "$from-1"
            }
          }
        },
        {
          "field name" : "TOTAL_BAGS",
          "field expression" :
          {
            "iterator kind" : "FUNC_SUM",
            "input iterator" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "TOTAL_BAGS",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$from-1"
              }
            }
          }
        }
      ]
}
```

Explanation:

- The root iterator of this query plan is a **RECEIVE** iterator with a single child (input iterator) that is a **SELECT** iterator.
- The value of the **`FROM`** field is an iterator; in this case, it is a **TABLE** iterator.
- The index **`compindex_namephone`** is used and the index is not covering (as you need to scan the table to fetch columns other than the index entries).
- Every value in the IN clause is evaluated using an index scan with an equality condition. There are four index scans that are performed each evaluating one equality condition.
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable

Copyright © 2022, 2025, Oracle and/or its affiliates.

(**$$bagdet**) is the same as the **row variable** of the TABLE iterator, as the index used is not covering.

- In the SELECT expression two fields (**fullname, bag.bagInfo[].tagNum**) are fetched. These correspond to two field names and field expressions in the SELECT expression clause. For the first field, the field expression is computed using **FIELD_STEP** iterator. For the second field, an **ARRAY_CONSTRUCTOR** iterator is used which iterates over the corresponding arrays to fetch the field value.

Query 9: A SINGLE PARTITION query using a primary index

Select the ticket details (ticket number, reservation code, tag number, and routing) for a passenger with a specific ticket number and reservation code.

```
SELECT fullName,bag.ticketNo, bag.confNo, bag.bagInfo[].tagNum,
bag.bagInfo[].routing FROM BaggageInfo bag WHERE
bag.ticketNo=1762311547917 AND bag.confNo="FH7G1W"
```

Plan:

```
{
  "iterator kind" : "RECEIVE",
  "distribution kind" : "SINGLE_PARTITION",
  "input iterator" :
  {
    "iterator kind" : "SELECT",
    "FROM" :
    {
      "iterator kind" : "TABLE",
      "target table" : "BaggageInfo",
      "row variable" : "$$bag",
      "index used" : "primary index",
      "covering index" : false,
      "index scans" : [
        {
          "equality conditions" : {"ticketNo":1762311547917},
          "range conditions" : {}
        }
      ]
    },
    "FROM variable" : "$$bag",
    "WHERE" :
    {
      "iterator kind" : "EQUAL",
      "left operand" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "confNo",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$$bag"
        }
      },
      "right operand" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
          {
            "iterator kind" : "CONST",
            "value" : "FH7G1W"
          }
        },
        "SELECT expressions" : [
          {
            "field name" : "fullName",
            "field expression" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "fullName",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$$bag"
              }
            }
          },
          {
            "field name" : "ticketNo",
            "field expression" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "ticketNo",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$$bag"
              }
            }
          },
          {
            "field name" : "confNo",
            "field expression" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "confNo",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$$bag"
              }
            }
          },
          {
            "field name" : "tagNum",
            "field expression" :
            {
              "iterator kind" : "ARRAY_CONSTRUCTOR",
              "conditional" : true,
              "input iterators" : [
                {
                  "iterator kind" : "FIELD_STEP",
                  "field name" : "tagNum",
                  "input iterator" :
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                {
                  "iterator kind" : "ARRAY_FILTER",
                  "input iterator" :
                  {
                    "iterator kind" : "FIELD_STEP",
                    "field name" : "bagInfo",
                    "input iterator" :
                    {
                      "iterator kind" : "VAR_REF",
                      "variable" : "$$bag"
                    }
                  }
                }
              }
            ]
          }
        },
        {
          "field name" : "routing",
          "field expression" :
          {
            "iterator kind" : "ARRAY_CONSTRUCTOR",
            "conditional" : true,
            "input iterators" : [
              {
                "iterator kind" : "FIELD_STEP",
                "field name" : "routing",
                "input iterator" :
                {
                  "iterator kind" : "ARRAY_FILTER",
                  "input iterator" :
                  {
                    "iterator kind" : "FIELD_STEP",
                    "field name" : "bagInfo",
                    "input iterator" :
                    {
                      "iterator kind" : "VAR_REF",
                      "variable" : "$$bag"
                    }
                  }
                }
              }
            ]
          }
        }
      ]
    }
  }
```

Explanation:

- The root iterator of this query plan is a RECEIVE iterator with a single child (input iterator) that is a SELECT iterator.
- This 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- The value of the **FROM** field is an iterator; in this case, it is a **TABLE** iterator.
- A **SINGLE_PARTITION** query can reference a primary index or a secondary index. The **primary key index** is used in this example. The index is not covering (as you need to scan the table to fetch columns other than the index entries).
- The index scan property contains the start and stop conditions that define the index scans to be performed.
- The FROM variable is the name of a variable ranging over the records produced by the FROM iterator. Here the FROM iterator is a TABLE iterator, and the FROM variable (**$$bag**) is the same as the **row variable** of the TABLE iterator, as the index used is not covering.
- In the SELECT expression five fields (**fullname, ticketNo,confNo, bag.bagInfo[].tagNum,bag.bagInfo[].routing**) are fetched. These correspond to five field names and field expressions in the SELECT expression clause. For the first three fields, the field expression is computed using **FIELD_STEP** iterator. For the last 2 fields, an **ARRAY_CONSTRUCTOR** iterator is used which iterates over the corresponding arrays to fetch the field value.

Query 10: Group data with fields not part of any index

Fetch the source of passenger bags and the count of bags for all passengers and group the data by the source.

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

Plan:

```
{
  "iterator kind" : "GROUP",
  "input variable" : "$gb-2",
  "input iterator" :
  {
    "iterator kind" : "RECEIVE",
    "distribution kind" : "ALL_PARTITIONS",
    "input iterator" :
    {
      "iterator kind" : "GROUP",
      "input variable" : "$gb-1",
      "input iterator" :
      {
        "iterator kind" : "SELECT",
        "FROM" :
        {
          "iterator kind" : "TABLE",
          "target table" : "BaggageInfo",
          "row variable" : "$bag",
          "index used" : "primary index",
          "covering index" : false,
          "index scans" : [
            {
              "equality conditions" : {},
              "range conditions" : {}
            }
          ]
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
      },
      "FROM variable" : "$bag",
      "FROM" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "fltRouteSrc",
        "input iterator" :
        {
          "iterator kind" : "ARRAY_SLICE",
          "low bound" : 0,
          "high bound" : 0,
          "input iterator" :
          {
            "iterator kind" : "FIELD_STEP",
            "field name" : "flightLegs",
            "input iterator" :
            {
              "iterator kind" : "FIELD_STEP",
              "field name" : "bagInfo",
              "input iterator" :
              {
                "iterator kind" : "VAR_REF",
                "variable" : "$bag"
              }
            }
          }
        }
      },
      "FROM variable" : "$flt_src",
      "SELECT expressions" : [
        {
          "field name" : "SOURCE",
          "field expression" :
          {
            "iterator kind" : "VAR_REF",
            "variable" : "$flt_src"
          }
        },
        {
          "field name" : "COUNT",
          "field expression" :
          {
            "iterator kind" : "CONST",
            "value" : 1
          }
        }
      ]
    },
    "grouping expressions" : [
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "SOURCE",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$gb-1"
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
          }
        }
      ],
      "aggregate functions" : [
        {
          "iterator kind" : "FUNC_COUNT_STAR"
        }
      ]
    }
  },
  "grouping expressions" : [
    {
      "iterator kind" : "FIELD_STEP",
      "field name" : "SOURCE",
      "input iterator" :
      {
        "iterator kind" : "VAR_REF",
        "variable" : "$gb-2"
      }
    }
  ],
  "aggregate functions" : [
    {
      "iterator kind" : "FUNC_SUM",
      "input iterator" :
      {
        "iterator kind" : "FIELD_STEP",
        "field name" : "COUNT",
        "input iterator" :
        {
          "iterator kind" : "VAR_REF",
          "variable" : "$gb-2"
        }
      }
    }
  ]
}
```

Explanation:

- In this query, you group passenger bags based on the flight source and determine the total number of bags belonging to one flight source.
- As the GROUP BY field (`bagInfo.flightLegs[0].fltRouteSrc` 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 `fltRouteSrc` and `count` of bags. 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 flight source 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 **primary key index** is used here and in this example, it is not a covering index as the query has fields that are not part of the entries of the primary index.
- 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. Every time a next() call on the FROM iterator returns true, the variable will be bound to the result produced by that iterator. In this example, the FROM variable is the row variable as the index is not covering.
- This row variable (**$bag**) will be referenced by iterators implementing the other clauses of the inner SELECT expression.
- The GROUP iterator creates an internal variable (`$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 `fltRouteSrc` groups and for the `fltRouteSrc` groups that it does contain, the count may be a partial sum (because all rows for a given `fltRouteSrc` 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 the 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 (**$gb-2**) that iterates over the records produced by the outer SELECT. The **$gb-2** variable has the `fltRouteSrc` and count of all bags grouped by `fltRouteSrc`.
- In the SELECT expression, two fields are fetched: `fltRouteSrc,count(*)`. These correspond to two field names and field expressions in the SELECT expression clause. For the first field, the field expression uses **FIELD_STEP** iterator. The second field is the aggregate function **count**. The iterator **FUNC_SUM** is used to iterate the result produced by its parent iterator and determine the total number of bags.

Table Modelling and Design

A critical part of the application development process is the task of modeling your data.

Proper modeling of your data is crucial to application performance, extensibility, application correctness, and finally, the ability for your application to support rich user experiences. In this article, you will learn some crucial aspects of data modeling and understand guidelines on how to model your persistent data for an Oracle NoSQL Database application.

The Oracle NoSQL Database gives the data modeler a large range of flexibility with respect to modeling application data. Understanding the tradeoffs associated with each level of flexibility is extremely useful in making wise data modeling decisions.

- Schema Flexibility in Oracle NoSQL Database
- Choice of Keys in NoSQL Database
- Using Indexes in NoSQL Database
- Transactions in NoSQL database

Copyright © 2022, 2025, Oracle and/or its affiliates.

Schema Flexibility in Oracle NoSQL Database

Unlike the relational database world with purely fixed schemas, NoSQL Database is largely about schema flexibility – that is the ability to easily change how data is organized and stored.

Schema flexibility in Oracle NoSQL Database mostly takes the form of non-scalar data types. These non-scalar data types can be used to embed flexible structures inside your tables.

Non-scalar data types:

Oracle NoSQL database supports the following non-scalar data types:

- **JSON** – JSON is a map of key/value pairs that can be used as a datatype of a column in the Oracle NoSQL Database. The JSON datatype gives you the ability to dynamically read and write attributes having no prior knowledge of what is and what is not stored in the JSON document. You can introspect into the document by reading from Oracle NoSQL Database as a JSON string, or you can specify path expressions as deep as you like into a hierarchy of JSON. As an example, you can create a JSON document that represents the variable terms and conditions of a contract. The document attribute names (or keys) can represent the tag or **name** of the contractual term or condition and the value of the attribute can represent the text of that term. Using a JSON brings you ultimate flexibility in your data model.
- **Records** – Records containing scalar or non-scalar values can be used as a datatype for a column in an Oracle NoSQL Database table. You can think of a record as a document with a fixed set of attributes, however noting that one or more attributes of the record can be a non-fixed array or JSON document, giving you the flexibility to extend a fixed document without modifying the schema. Records present an interesting intermediate step between the benefits of the fixed schema world (single copy of a schema) and the ultimate flexibility of the JSON world.
- **Arrays** – Arrays of scalar or non-scalar values can be used as a datatype for a column in an Oracle NoSQL Database table. Arrays can be convenient for storing a collection of event values. For example, you may wish to collect a list of behavioral segments for users as they browse web pages.

Trade-offs while using Flexible Schema:

Some guidelines that you can follow while considering flexible schema are listed below.

The Flexibility/Cost of scale Tradeoff :

When thinking about how flexible you want your schema to be, it’s important to understand that the more flexible you make your schema, the bigger the challenge is for scaling your solution. For example, let’s say that you are storing information on user behavior. And you want to store this information as the users access your website. You can implement one of the two options here. You can choose to model the solution using fixed columns for the required user attributes that you will need to track. Alternatively, you can choose to model this using a JSON document, giving you the flexibility to add and remove attributes for users without having to evolve your schema.

The second option may work quite well for small numbers of users; however, if you will need to scale this solution to large numbers of users, then you need extra storage. You also need additional compute overhead for processing the key/value pairs (attribute names and their values) in the JSON document. This could make the cost of scaling your solution prohibitive. Extra storage is needed to store the metadata along with the data (e.g. the attribute names) and extra compute is needed to serialize and de-serialize these documents. If you are using a replication factor of more than one, then that adds additional overhead for each tracked user. If

Copyright © 2022, 2025, Oracle and/or its affiliates.

a large scale is a major requirement for you then you'll most likely want to trade off flexibility for storage efficiency and consider using a more fixed schema.

The Flexibility/Latency Tradeoff :

In many NoSQL applications, low latency data access is a key requirement. In these situations, it's important to understand the potential tradeoff with respect to the I/O latency of using one data modeling method over another. In this respect, using a non-scalar data type such as a record, array, or JSON document will entail a read followed by an update.

For example, when you add a new value to an array your application must read the record from NoSQL Database first, add the value to the array, and then write it back to NoSQL Database. Even if performing this operation using the SQL UPDATE operator of Oracle NoSQL Database (which executes in the replication node), the record must still be read from persistent storage, de-serialized, modified, serialized, and written back. On a system with a spinning disk, this could cost anywhere from fifteen to thirty milliseconds (or more). For certain applications like online advertising, this may be beyond the latency SLA that can be tolerated. If you are faced with similar stringent latency SLAs then you should consider favoring a child table approach which eliminates the read and will allow you to simply perform a write of the new value. Of course, the tradeoff here is one of flexibility for low latency.

For more information on when to use parent-child tables, see Using Parent-Child tables in Oracle NoSQL Database.

Updates to Non-Scalars versus Inserts :

The Oracle NoSQL Database storage engine is based on an append-only architecture, also known as log-structured storage. Log structured storage systems perform extremely well for insert operations where the new records to be inserted are simply appended to the end of the log. Update operations involve appending the updated record to the log and then marking the old record for deletion. Records marked for deletion are regularly cleaned from NoSQL Database's logs to free up disk space by a background process called the cleaner. Although the cleaner is highly optimized, it will add some CPU and I/O overhead to the replication node. The more updates performed by your application, the more log cleaning activity there will be.

As a guideline, if you have extreme performance goals for your application (or for a specific table), you should strongly consider trying to craft your data model by using parent/child tables versus non-scalar columns, giving you the potential for replacing updates with inserts. For more information on when to use parent-child tables, see Using Parent-Child tables in Oracle NoSQL Database.

Static Vs Dynamic Data :

In many applications, it's possible to identify portions of the data that are somewhat static and change relatively slowly, and other portions of the data which are highly dynamic and change frequently, even at millisecond granularity. For example, in online advertising, campaigns are a relatively slow-moving piece of data while the budget spent (impressions or clicks delivered) can change every few milliseconds as millions of users load web pages that have ads associated with the campaign. The data pertaining to budgets is a case of highly dynamic data. This is an example of a scenario that has a high velocity of write operations. Oracle NoSQL database is a log-structured, append-only storage architecture, where inserts are more optimal than an update operation.

For the more static portions of your data, the flexibility of the non-scalar datatypes may be an attractive option for your application. Using a JSON document could provide an extensible way for your application to interact with this data without undue sacrifice to performance. On the flip side, for data that is changing rapidly or being inserted rapidly, you'll want to consider trading off flexibility for this data and use a parent table with a fixed schema and a child table with a fixed schema. Whether or not you choose to model the rapidly changing data as a parent or

Copyright © 2022, 2025, Oracle and/or its affiliates.

child table will depend largely on how you wish to access it. For more information on when to use parent-child tables, see Using Parent-Child tables in Oracle NoSQL Database.

JSON Collection Tables:

Oracle NoSQL Database supports JSON Collection tables. This is particularly useful for applications that store and retrieve their data purely as documents. Such tables contain primary key fields and a document. The schema of the JSON Collection table can't be altered to add typed fields. JSON Collection tables are created to simplify the management and manipulation of documents. The JSON collection table eliminates the need to declare fields as type JSON during table creation. When you insert data into the table, each row is inserted as a single document containing any number of JSON fields. You can add JSON fields through INSERT or UPSERT operations.

Choice of Keys in NoSQL Database

Primary keys and shard keys are important elements in your schema and help you access and distribute data efficiently.

Primary keys and shard keys are indispensable for data distribution and easy accessibility. You specify primary keys and shard keys only when you create a table. They remain in place for the life of the table, and cannot be changed or dropped.

Using Primary Keys and Shard Keys in Oracle NoSQL tables

Primary Keys

You must designate one or more primary key columns when you create your table. The primary key cannot be changed and exists for the life of the table. A primary key uniquely identifies every row in the table. For simple CRUD operations, Oracle NoSQL Database uses the primary key to retrieve a specific row to read or modify. Since the underlying storage in NoSQL Database is based on a key/value model, the choice of the primary key can greatly enhance the performance of certain lookup operations.

Shard Keys

The main purpose of shard keys is to distribute data across the Oracle NoSQL Database cluster for scalability and to co-locate the records that share the same shard key on the same physical node for easy reference. These records can be accessed atomically and efficiently.

Impact of keys while developing an application:

In an Oracle NoSQL Database, replication nodes are grouped together to form the shards of the NoSQL Database cluster. When an application asks to retrieve the record for a given key, the NoSQL Database driver will hash a portion of the key (denoted as the **shard key**) to identify the shard that houses the data. Once the shard is identified, the NoSQL Database driver can choose to read the data from the most optimal replica in the shard, depending on the requested consistency level. With respect to the write operations, the NoSQL Database driver will always route the write requests to the dynamically elected leader node of the shard. Hence, from the perspective of workload scaling, you can generally think of this architecture as being scaled by adding shards. Oracle NoSQL Database supports the online elastic expansion of the cluster by adding shards, however, without the proper selection of a shard key, expanding the cluster will be useless in scaling your solution.

How you design primary keys and shard keys has huge implications on scaling and realizing the system throughput. For instance, when records share shard keys, you can delete multiple table rows in an atomic operation, or retrieve a subset of rows in your table in a single atomic operation. In addition to enabling scalability, well-designed shard keys can improve performance by requiring fewer cycles to put data on, or get data from, a single shard. Shard

Copyright © 2022, 2025, Oracle and/or its affiliates.

keys designate storage on the same shard to facilitate efficient queries for key values. However, because you want your data to be distributed across the shards for best performance and scalability, you will want to avoid shard keys that have a small number of unique values.

Important factors to consider when choosing a shard key:

- **Cardinality:** Low cardinality field groups are stored together on a small number of shards. In turn, those shards require frequent data rebalancing, increasing the likelihood of hot shard issues. Instead, each shard key should have high cardinality, where it can express several million values. For best performance and value, choose fields with high cardinalities, such as identity numbers, where millions of records are possible.
- **Atomicity:**Only objects that share the same shard key can participate in a transaction. If you have a requirement for ACID transactions that span multiple records, choose only a shard key that lets you meet that requirement.

Best practices to follow:

- **Uniform distribution of shard keys:**Operations may be limited by the capacity of a single shard. When shard keys are uniformly distributed, no single shard limits the capacity of the system. Choosing one or more columns whose values are known to be uniformly distributed is ideal.
- **Query Isolation:** Queries should be targeted to a specific shard to maximize scalability. If queries are not isolated to a single shard, the query will be applied to all shards. This is less efficient and increases query latency. Make sure your queries fetch data stored in a single shard. Well-designed shard keys can improve performance by getting data from a single shard. Shard keys designate storage on the same shard to facilitate efficient queries for key values. Specify the fields (which are frequently used in your application queries) as shard keys.

Key Sizes and Key Only Modeling Methods

Oracle NoSQL Database caches the keys for each table. So the key size is a critical component to the effective use of memory and ultimately may be a determining factor in the ability of Oracle NoSQL Database to service your performance SLAs. Hence, it is important for you to create primary keys that are as efficient as possible with respect to size. For workloads that require very low latencies for the read and writes (single to low double-digit milliseconds) across millions of operations per second, exploiting cached keys in NoSQL's B-trees can be the make or break of building an application capable of achieving these stringent requirements. Furthermore, if you can encode what would otherwise be non-key values as part of the primary key and also size your keys and the NoSQL Database cluster carefully, then you can realize the enormous benefits of memory cached B-tree access methods that are maintained with ACID semantics. For highly optimal, ultra-low latency applications, Oracle NoSQL provides key-only accessors for those workloads that can model everything as key-only data. Oracle NoSQL Database offers convenient key-only access APIs such as `multiKetGeys` and `tableKeysIterator` for doing key-only scans.

When considering whether or not key-only modeling of your data is right for your application, you should consider the following:

- **Latency and throughput SLAs** – Do you have very stringent latency and throughput SLAs that would require a key-only model? Can you afford to perform an I/O when retrieving a value, noting that for spinning disks, the average latency of retrieving your value could be anywhere from fifteen to thirty milliseconds and for Single Shared Disk (SSD) this could be anywhere from one to 5 milliseconds.
- **Spinning disk versus SSDs** – If you are considering using SSDs and your latency SLAs are for reads that can comfortably fit within the 5-millisecond range then it's probably not worth the effort to try and craft a key-only model for your application.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- **Code maintainability and extensibility** – Key-only modeling brings large performance benefits to your application at the potential cost of code maintainability and extensibility. You may find that encoding your value into the key can ultimately be a complex and esoteric strategy. Ultimately, you will have to make a judgment call on whether or not the code you develop and maintain is too complex and esoteric to be worth the benefit of the key-only solution.
- **Accurate sizing data** – Is it possible for you to derive a somewhat accurate sizing of your keys such that you can adequately size the Oracle NoSQL Database cluster? Sizing the cluster and the cache of each replication node will be crucial to exploiting the benefits of a key-only data model.

Key Column Ordering and Query-ability

In Oracle NoSQL Database, the order of declaration for key columns is crucial to satisfying partial key lookup queries. This is because of the way that the storage engine manages the underlying B-trees. You can think of composite keys as an ordered concatenation of the columns specified in the DDL for the key declaration (primary key or index key). You should think of the order from the most significant column to the least significant column based on the appearance of the columns in the DDL for the key. If your table has a composite primary key (a primary key with more than one column), then the primary key becomes a concatenation of the string representation of each column. Here for better performance of queries, it is important to specify the most commonly used query column as the most significant column in the primary key.

As you start to think about how you will size your cluster and your Oracle NoSQL Database caches, a critical consideration is to get an estimate of your key sizes. Sizing your caches so that Oracle NoSQL can keep most or all of the index nodes in memory can help your application realize enormous performance benefits. Understanding how keys are serialized and stored persistently can help you in getting a more accurate sizing estimate. In Oracle NoSQL Database, numeric keys are stored as compressed String values but must remain sortable when in string format. This means that a numeric key must be a fixed size when represented as a key string. See Initial Capacity Planning for more details on shard capacity, shard storage, and throughput capacities and how to estimate total shards and machines.

Choice of using Identity column Vs UUID

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 sequence generator 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.

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 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. 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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
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.
Identity columns allow Oracle NoSQL to automatically assign values within a single region. An error is thrown if an Identity column is used in a multi-region deployment.	UUID columns allow Oracle NoSQL to automatically assign global values across regions. This is useful in multi-region deployments. A UUID column is larger than an Identity column.

Using Indexes in NoSQL Database

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 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.

In an 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?

There are no statistics on the number and distribution of values in a table column. 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. You can use index hints to force the use of a particular index in queries. You can use a query execution plan to understand what indexes are being used in the query. For more information on how a query is executed, see Query execution plan.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Secondary Index

There will be cases where you will want to use a secondary index to support some of your read requirements. Each secondary index that you add to a table will incur some overhead for writes as each index will need to be maintained. The good news with Oracle NoSQL is that secondary index partitions live on the same shard as the primary data, so the updates to the secondary index are limited on a per-shard basis. Index updates in Oracle NoSQL are also atomic, so your application can be guaranteed that updates to records in the shard are consistent with updates to the secondary index and these structures will never be out of sync. Another factor for consideration is that Oracle NoSQL Database nodes will keep the non-leaf index nodes in the cache, and will never cache the leaf portion (i.e. the data record). This gives the indexed scan an enormous performance benefit (for systems using spinning disk) over the non-indexed scan.

There are several things that you should think about when deciding on using a secondary index in Oracle NoSQL Database:

- Filtering data close to the source – In Oracle NoSQL Database, secondary indexes are the primary mechanism for you to utilize when your query needs a filter and that filter needs to be executed as close as possible to the data. To fully understand why you may need a secondary index to filter your data for querying, let's consider your options for scanning the data in a table:
  - Unordered parallel table scan with no full shard key – The shard key is a table column or multiple columns used to control how the rows of that table are distributed. The main purpose of shard keys is to distribute data across the Oracle NoSQL Database Cloud cluster for scalability, and to position records that share the same shard key locally for easy reference and access. When you write a query using filters as columns that are part of the shard key but also include other columns, then you end up doing a parallel table scan. Each shard is scanned in parallel and the data is returned to your application. This will return every record in the table across all shards in the NoSQL Database.
  - Ordered or unordered parallel index scan – The B-tree index at each shard is scanned in parallel. If an ordered scan is requested, the results are merged and presented.
- Each option for scanning a table has its own costs and benefits and you should carefully weigh these tradeoffs and use what you know about the application requirements and expected workload to help guide your modeling decision.
  - Efficient range scans – Will it be common for your queries to restrict the value ranges? For example, if your application needs to answer queries like "find all records between a range of dates" then using secondary indexes in Oracle NoSQL Database will be the easiest and most efficient way for your application to answer these types of queries.
  - Workload and index maintenance update – Is it acceptable for writes to incur some extra overhead for index maintenance? Does your workload exhibit heavy read activity where latency for reads is more important than incurring extra write overhead?

See Tuning and Optimizing SQL queries for more guidelines on using indexes in queries.

Transactions in NoSQL database

In Oracle NoSQL Database, a transaction is treated as a logical, atomic unit of work that entails a single database operation.

Every data modification in the database takes place in a single transaction, managed by the system. Database developers do not have the ability to group multiple operations into a single

Copyright © 2022, 2025, Oracle and/or its affiliates.

transaction because there isn't the notion of begin/end transactions. In a database, transactional semantics are often described in terms of ACID properties.

ACID properties

In Oracle NoSQL Database, transactions maintain all the following properties and developers can control some of them.

- **Atomicity:** Transaction either completes or fails in its entirety. There is no in-between state or no partial transactions.
- **Consistency:** Transaction leaves the database in a valid state.
- **Isolation:** No two transactions mingle or interfere with each other. Developers get the same result when the two transactions are executed in sequence or executed in parallel.
- **Durability:** Changes in a transaction are saved and the changes survive any type of failure (network, disk, CPU, or a power failure).

Developers can define a wide range of consistency levels depending on the application's needs with the Oracle NoSQL Database Direct Driver. In addition, the Oracle NoSQL Database Drivers (commonly called the SDKs) support eventual and absolute consistency.

Developers can also configure durability such that updated rows in the database survive any failure with the Oracle NoSQL Database Direct Driver. Durability is not configurable in the SDKs.

Atomicity and Isolation are not configurable but Oracle NoSQL Database allows you to control consistency and durability policies in order to trade-off the performance for application needs. Some NoSQL databases only support eventual consistency but have no mechanism for absolute consistency.

Shard keys play an important role in achieving the ACID properties in the Oracle NoSQL database. For instance, when records share shard keys, you can delete multiple table rows in an atomic operation, or retrieve a subset of rows in your table in a single atomic operation. In addition to enabling scalability, well-designed shard keys can improve performance by requiring fewer cycles to put data on, or get data from, a single shard.

The NoSQL table hierarchy is an ideal data model for applications that need some data normalization, but also require predictable, low latency at scale. The hierarchy links different tables to enable left outer joins, combining rows from two or more tables based on related columns between them. Such joins execute efficiently since rows from the parent-child tables are co-located in the same shards. Also, writes to multiple tables in a table hierarchy obey transactional ACID properties since the records residing in each table of the hierarchy share the same shard key. All write operations perform as a single atomic unit. So all of the write operations will execute successfully, or none of them will.

Using Parent-Child tables in the Oracle NoSQL database

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

Many NoSQL databases support data types like arrays and maps. When modeling a data relationship, application developers may find it easier to have each parent row store its child rows inside an array or a map in a nested structure. By doing so, not only is the data relationship denormalized but it has the potential for creating large parent rows, especially when the hierarchy is heavily nested, resulting in inefficient storage and poor performance. Oracle NoSQL Database’s table hierarchy is the ideal data model to avoid issues associated with arrays and maps. One of the biggest benefits of using child tables over embedded arrays is for those workloads that have a high velocity of write operations. When using embedded arrays, the write operations become updates, but when they are modeled as child tables, those

Copyright © 2022, 2025, Oracle and/or its affiliates.

operations become inserts. Inserts in a log-structured, append-only storage architecture are much more optimal than updates. Utilizing a table hierarchy should be considered when building data relationships in Oracle NoSQL Database.

The NoSQL table hierarchy is an ideal data model for applications that need some data normalization, but also require predictable, low latency at scale. The hierarchy links different tables to enable left outer joins, combining rows from two or more tables based on related columns between them. Such joins execute efficiently since rows from the parent-child tables are co-located in the same shards. Also, writes to multiple tables in a table hierarchy obey transactional ACID properties since the records residing in each table of the hierarchy share the same shard key. All write operations perform as a single atomic unit. So all of the write operations will execute successfully, or none of them will.

The Benefits of a Table Hierarchy

Oracle NoSQL Database table hierarchy comes with the following benefits:

- **Highly efficient for storing data in a parent-child hierarchy** - Parent and child rows are stored in separate NoSQL tables, reducing the size of parent rows compared with the single parent with child rows in nested arrays or maps. Write operations on parent or child tables create new versions of smaller rows and store these changes efficiently, given the append-only architecture of Oracle NoSQL Database.
- **Highly performant for read and write workloads** - Parent and child rows reside in the same local shard, enabling write and read operations to achieve high performance since all records in the hierarchy can be read or written in a single network call.
- **Highly flexible for fine-grained authorization** - Access rights to a parent or child table can be configured individually based on conditions at run-time, offering granular and flexible authorization.
- **Scalable ACID transactions** - Uniquely balance the goals of scalability, low latency, and ACID by co-locating parent and child data on the same shard.
- **Table joins** - Data can be queried using the nested table clause or left outer joins.

Characteristics of parent-child tables:

- A child table inherits the primary key columns of its parent table.
- All tables in the hierarchy have the same shard key columns, which are specified in the create table statement of the root table.
- A parent table cannot be dropped before its children are dropped.
- A referential integrity constraint is not enforced in a parent-child table.

A NoSQL table hierarchy not only captures the relationship between data entities but also takes advantage of the co-location of the parent-child rows to offer highly performant retrievals and superior scalability. The table hierarchy enables applications to implement ACID transactions. All data in the same parent-child rows are stored in the same shard and can be committed as a single database operation to ensure atomicity, consistency, isolation, durability.

Handling Errors

Learn how to handle errors and exceptions.

- Handling Driver Errors

Copyright © 2022, 2025, Oracle and/or its affiliates.

Handling Driver Errors

Learn how to handle driver-related errors and exceptions reported during building or running the application.

- Java
- Python
- Go
- Node.js
- C#

Java

Java errors are thrown as exceptions when you build or run your application. The `NoSQLException` class is the base for most exceptions thrown by the driver. However, the driver throws exceptions directly for some classes, such as `IllegalArgumentException` and `NullPointerException`.

In general, NoSQL exception instances are split into two broad categories:

- Exceptions that may be retried with the expectation that they may succeed on retry.

  These exceptions are instances of the `RetryableException` class. These exceptions usually indicate resource consumption violations.
- Exceptions that will fail even after retry.

  Examples of exceptions that should not be retried are `IllegalArgumentException`, `TableNotFoundException`, and any other exception indicating a syntactic or semantic error.

Python

Python errors are raised as exceptions defined as part of the API. They are all instances of Python's RuntimeError. Most exceptions are instances of `borneo.NoSQLException` which is a base class for exceptions raised by the Python driver.

Exceptions are split into 2 broad categories: Exceptions that may be retried with the expectation that they may succeed on retry. These are all instances of `borneo.RetryableException`. Examples of these are the instances of `borneo.ThrottlingException` which is raised when resource consumption limits are exceeded. Exceptions that should not be retried, as they will fail again. Examples of these include `borneo.IllegalArgumentException` , `borneo.TableNotFoundException`, etc.

`borneo.ThrottlingException` instances will never be thrown in an on-premise configuration as there are no relevant limits.

Go

Go SDK errors are reported as `nosqlerr.Error` values defined as part of the API. Errors are split into 2 broad categories:

- Errors that may be retried with the expectation that they may succeed on retry. These are retryable errors on which the `Error.Retryable()` method call returns `true`. Examples of

Copyright © 2022, 2025, Oracle and/or its affiliates.

these include `nosqlerr.OperationLimitExceeded`, `nosqlerr.ReadLimitExceeded`, `nosqlerr.WriteLimitExceeded`, which are raised when resource consumption limits are exceeded.

- Errors that should not be retried, as they will fail again. Examples of these include `nosqlerr.IllegalArgumentError`, `nosqlerr.TableNotFoundError`, etc.

Node.js

Asynchronous methods of `NoSQLClient` return Promise as a result and if an error occurs it results in the Promise rejection with that error. For synchronous methods such as `NoSQLClient` constructor errors are thrown as exceptions. All errors used by the SDK are instances of `NoSQLError` or one of its subclasses. In addition to the error message, each error has `errorCode` property set to one of standard error codes defined by the ErrorCode enumeration. `errorCode` may be useful to execute conditional logic depending on the nature of the error.

For some error codes, specific subclasses of `NoSQLError` are defined, such as `NoSQLArgumentError`, `NoSQLProtocolError`, `NoSQLTimeoutError`, etc. `NoSQLAuthorizationError` may have one of several error codes depending on the cause of authorization failure. In addition, errors may have `cause` property set to the underlying error that caused the current error. Note that the cause is optional and may be an instance of an error that is not part of the SDK.

In addition, error codes are split into 2 broad categories:

- Errors that may be retried with the expectation that the operation may succeed on retry. Examples of these are `ErrorCode.READ_LIMIT_EXCEEDED` and `ErrorCode.WRITE_LIMIT_EXCEEDED` which are throttling errors (relevant for the Cloud environment), and also `ErrorCode.NETWORK_ERROR` since most network conditions are temporary.
- Errors that should not be retried, as the operation will most likely fail again. Examples of these include `ErrorCode.ILLEGAL_ARGUMENT` (represented by `NoSQLArgumentError`), `ErrorCode.TABLE_NOT_FOUND`, etc.

You can determine if the `NoSQLError` is retryable by checking `retryable` property. Its value is set to true for retryable errors and is false or undefined for non-retryable errors.

Retry Handler

The driver will automatically retry operations on a retryable error. Retry handler determines:

- Whether and how many times the operation will be retried.
- How long to wait before each retry.

RetryHandler is an interface with with 2 properties:

- `RetryHandler#doRetry` that determines whether the operation should be retried based on the operation, number of retries happened so far and the error occurred. This property is usually a function, but may be also be set to boolean false to disable automatic retries.
- `RetryHandler#delay` that determines how long to wait before each successive retry based on the same information as provided to `RetryHandler#doRetry`. This property is usually a function, but may also be set to number of milliseconds for constant delay.

C#

NoSQLException serves as a base class for many exceptions thrown by the driver. However, in certain cases the driver uses standard exception types such as:

Copyright © 2022, 2025, Oracle and/or its affiliates.

- `ArgumentException` and its subclasses such as `ArgumentNullException`. They are thrown when an invalid argument is passed to a method or when an invalid configuration (in code or in JSON) is passed to create NoSQLClient instance.
- `TimeoutException` is thrown when an operation issued by `NoSQLClient` has timed out. If you are getting many timeout exceptions, you may try to increase the timeout values in `NoSQLConfig` or in options argument passed to the `NoSQLClient` method.
- `InvalidOperationException` is thrown when the service is an invalid state to perform an operation. It may also be thrown if the query has failed be cause its processing exceeded the memory limit specifed in `QueryOptions.MaxMemoryMB` or `NoSQLConfig.MaxMemoryMB`. In this case, you may increase the corresponding memory limit. Otherwise, you may retry the operation.
- `InvalidCastException` and `OverflowException` may occur when working with sublcasses of FieldValue and trying to cast a value to a type it doesn't support or cast a numeric value to a smaller type causing arithmetic overflow.
- `OperationCanceledException` and `TaskCanceledException` if you issued a cancellation of the operation started by a method of `NoSQLClient` using the provided CancellationToken.

In addition, exceptions may be split into two broad categories:

- Exceptions that may be retried with the expectation that the operation may succeed on retry. In general these are subclasses of `RetryableException`. These include throttling exceptions as well as other exceptions where a resource is temporarily unavailable. Some other subclasses of `NoSQLException` may also be retryable depending on the conditions under which the exception occurred. In addition, network-related errors are retryable because most network conditions are temporary.
- Exceptions that should not be retried because they will still fail after retry. They include exceptions such as `TableNotFoundException`, `TableExistsException` and others as well as standard exceptions such as `ArgumentException`.

You can determine if a given instance of `NoSQLException` is retryable by checking its `IsRetryable` property.

Retry Handler

By default, the driver will automatically retry operations that threw a retryable exception (see above). The driver uses retry handler to control operation retries. The retry handler determines:

- Whether and how many times the operation will be retried.
- How long to wait before each retry.

All retry handlers implement `IRetryHandler` interface. This interface provides two methods, one to determine if the operation in its current state should be retried and another to determine a retry delay before the next retry. You have a choice to use default retry handler or set your own retry handler as `RetryHandler` property of `NoSQLConfig` when creating `NoSQLClient` instance.

Note

Retries are only performed within the timeout period alloted to the operation and configured as one of timeout properties in `NoSQLConfig` or in options passed to the `NoSQLClient` method. If the operation or its retries have not succeded before the timeout is reached, `TimeoutException` is thrown.

Copyright © 2022, 2025, Oracle and/or its affiliates.

By default, the driver uses `NoSQLRetryHandler` class which controls retires based on operation type, exception type and whether the number of retries performed has reached a preconfigured maximum. It also uses exponential backoff delay to wait between retries starting with a pre configured base delay. You may customize the properties such as maximum number of retries, base delay and others by creating your own instance of `NoSQLRetryHandler` and setting it as a `RetryHandler` property in `NoSQLConfig`. For example:

```
var client = new NoSQLClient(
    new NoSQLConfig
    {
        Region = .....,
        ...............
        RetryHandler = new NoSQLRetryHandler
        {
            MaxRetryAttempts = 20,
            BaseDelay = TimeSpan.FromSeconds(2)
        }
    });
```

If you don't specify the retry handler, the driver will use an instance of `NoSQLRetryHandler` with default values for all parameters. Alternatively, you may choose to create your own retry handler class by implementing `IRetryHandler` interface. The last option is to disable retries all together. You may do this if you plan to retry the operations within your application instead. To disable retries, set `RetryHandler` property of `NoSQLConfig` to `NoRetries`:

```
var client = new NoSQLClient(
    new NoSQLConfig
    {
        Region = .....,
        ...............
        RetryHandler = NoSQLConfig.NoRetries
    });
```

Handle Resource Limits: Programming in a resource-limited environment can be challenging. Tables have user-specified throughput limits and if an application exceeds those limits it may be throttled, which means an operation may fail with one of the throttling exceptions such as `ReadThrottlingException` or `WriteThrottlingException`. This is most common when using queries, which can read a lot of data, using up capacity very quickly. It can also happen for get and put operations that run in a tight loop.

Even though throttling errors will be retried and using custom `RetryHandler` may allow more direct control over retries, an application should not rely on retries to handle throttling as this will result in poor performance and inability to use all of the throughput available for the table. The better approach would be to avoid throttling entirely by rate-limiting your application. In this context rate-limiting means keeping operation rates under the limits for the table.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Glossary

Copyright © 2022, 2025, Oracle and/or its affiliates.

Index

Copyright © 2022, 2025, Oracle and/or its affiliates.