

Oracle® NoSQL Database

Getting Started with SQL for Oracle NoSQL Database



Release 12.2.4.5

E85380-01

February 2018

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

ORACLE®

Copyright © 2017, 2018, Oracle and/or its affiliates. All rights reserved.

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

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

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Contents

Preface

Conventions Used in This Book

viii

1 Introduction to SQL for Oracle NoSQL Database

Part I Introductory Examples

2 Simple SELECT Queries

SQLBasicExamples Script	2-1
Running the SQL Shell	2-2
Choosing column data	2-2
Substituting column names for a query	2-3
Computing values for new columns	2-3
Identifying tables and their columns	2-4
Filtering results	2-5
Ordering Results	2-6
Limiting and Offsetting Results	2-8
Using External Variables	2-9

3 Working with complex data

SQLAdvancedExamples Script	3-1
Working with Timestamps	3-4
Working With Arrays	3-5
Working with Records	3-9
Using ORDER BY to Sort Results	3-11
Working With Maps	3-12
Using the size() Function	3-14

4 Working with JSON

SQLJSONExamples Script	4-1
Basic Queries	4-4
Using WHERE EXISTS with JSON	4-5
Seeking NULLS in Arrays	4-6
Examining Data Types JSON Columns	4-7
Using Map Steps with JSON Data	4-10
Casting Datatypes	4-11
Using Searched Case	4-12

5 Working With Indexes

Basic Indexing	5-1
Using Index Hints	5-2
Complex Indexes	5-3
Multi-Key Indexes	5-4
Indexing JSON Data	5-8

6 Modifying Table Rows using UPDATE Statements

Example Data	6-1
Changing Field Values	6-1
Modifying Array Values	6-3
Adding Elements to an Array	6-3
Changing an Existing Element in an Array	6-5
Removing Elements from Arrays	6-6
Modifying Map Values	6-8
Removing Elements from a Map	6-9
Adding Elements to a Map	6-9
Updating Existing Map Elements	6-12
Managing Time to Live Values	6-15
Avoiding the Read-Modify-Write Cycle	6-17

Part II Language Definition

7 The SQL for Oracle NoSQL Database Data Model

Example Data	7-1
Data Types and Values	7-2
Wildcard Types and JSON Data	7-3

JSON Data	7-4
Timestamp	7-4
Timestamp Functions	7-5
Type Hierarchy	7-6
Subtype-Substitution Rule Exceptions	7-8
SQL for Oracle NoSQL Database Sequences	7-8
Sequence Concatenation Function	7-9

8 SQL for Oracle NoSQL Database Queries

Select-From-Where (SFW) Expressions	8-1
SELECT Clause	8-2
SELECT Clause Hints	8-2
FROM Clause	8-3
WHERE Clause	8-4
ORDER BY Clause	8-4
Comparison Rules	8-5
OFFSET Clause	8-5
LIMIT Clause	8-5

9 Expressions

Path Expressions	9-1
Field Step Expressions	9-2
Map Filter Step Expressions	9-2
Array Filter Step Expressions	9-3
Array Slice Step Expressions	9-4
Constant Expressions	9-5
Column Reference Expression	9-5
Variable Reference Expression	9-5
Searched Case Expressions	9-6
Cast Expressions	9-6

10 Operators

Logical Operators	10-1
Value Comparison Operators	10-1
Sequence Comparison Operators	10-3
IS NULL Operator	10-3
Exists Operator	10-4

	Is-Of-Type Operator	10-4
11	Constructors	
	Array Constructors	11-1
	Map Constructors	11-1
12	Built-in Functions	
	Time to Live Functions	12-1
	Time Functions	12-2
13	SQL UPDATE Statements	
	Update Statement Syntax	13-1
	Update Clauses	13-2
	SET Clause	13-2
	ADD Clause	13-3
	PUT Clause	13-3
	REMOVE Clause	13-4
	SET TTL Clause	13-4
A	Introduction to the SQL for Oracle NoSQL Database Shell	
	Running the shell	A-1
	Configuring the shell	A-2
	Shell Utility Commands	A-2
	connect	A-3
	consistency	A-3
	describe	A-3
	durability	A-3
	exit	A-3
	help	A-4
	history	A-4
	import	A-4
	load	A-4
	mode	A-5
	output	A-8
	page	A-8
	show faults	A-8
	show query	A-8
	show tables	A-9

show users	A-9
show roles	A-9
timeout	A-9
timer	A-9
verbose	A-10
version	A-10

Preface

This document is intended to provide a rapid introduction to the SQL for Oracle NoSQL Database and related concepts. SQL for Oracle NoSQL Database is an easy to use SQL-like language that supports read-only queries and data definition (DDL) statements. This document focuses on the query part of the language. For a more detailed description of the language (both DDL and query statements) see the *SQL for Oracle NoSQL Database Specification*.

This book is aimed at developers who are looking to manipulate Oracle NoSQL Database data using a SQL-like query language. Knowledge of standard SQL is not required but it does allow you to easily learn SQL for Oracle NoSQL Database.

Conventions Used in This Book

The following typographical conventions are used within this manual:

Information that you are to type literally is presented in `monospaced font`.

Variable or non-literal text is presented in *italics*. For example: "Go to your *KVHOME* directory."

Case-insensitive keywords, like SELECT, FROM, WHERE, ORDER BY, are presented in UPPERCASE.

Case sensitive keywords, like the function size(item) are presented in lowercase.

**Note:**

Finally, notes of special interest are represented using a note block such as this.

1

Introduction to SQL for Oracle NoSQL Database

Welcome to SQL for Oracle NoSQL Database. This language provides a SQL-like interface to Oracle NoSQL Database that can be used from a command line interface, scripts, or from the Oracle NoSQL Database Java Table Driver. The SQL for Oracle NoSQL Database data model supports flat relational data, hierarchical typed (schema-full) data, and schema-less JSON data. SQL for Oracle NoSQL Database is designed to handle all such data in a seamless fashion without any "impedance mismatch" among the different sub models.

For information on the command line shell you can use to run SQL for Oracle NoSQL Database queries, see [Introduction to the SQL for Oracle NoSQL Database Shell](#). For information on executing SQL for Oracle NoSQL Database queries from the Oracle NoSQL Database Java Table Driver, see the *Oracle NoSQL Database Getting Started with the Table API* manual.

This book is broken into two parts:

- Part I provides an examples-based introduction to the language. It begins here: [Introductory Examples](#)
- Part II provides a textual description of the language. It begins here: [Language Definition](#)

Part I

Introductory Examples

This part provides an examples-based introduction to SQL for Oracle NoSQL Database. For a textual description of the language, please see [Language Definition](#).

2

Simple SELECT Queries

This section presents examples of simple queries for relational data. To follow along with the examples, get the `Examples` download from here and run the `SQLBasicExamples` script found in the `sql` folder. The script creates the table as shown, and imports the data.

SQLBasicExamples Script

The script `SQLBasicExamples` creates the following table:

```
CREATE TABLE Users (  
  id integer,  
  firstname string,  
  lastname string,  
  age integer,  
  income integer,  
  primary key (id)  
);
```

The script also load data into the `Users` table with the following rows (shown here in JSON format):

```
{  
  "id":1,  
  "firstname":"David",  
  "lastname":"Morrison",  
  "age":25,  
  "income":100000,  
}  
  
{  
  "id":2,  
  "firstname":"John",  
  "lastname":"Anderson",  
  "age":35,  
  "income":100000,  
}  
  
{  
  "id":3,  
  "firstname":"John",  
  "lastname":"Morgan",  
  "age":38,  
  "income":null,  
}  
  
{  
  "id":4,  
  "firstname":"Peter",  
  "lastname":"Smith",  
  "age":38,  
  "income":80000,  
}
```

```

}

{
  "id":5,
  "firstname":"Dana",
  "lastname":"Scully",
  "age":47,
  "income":400000,
}

```

You run the SQLBasicExamples script using the `load` command:

```

> cd <installdir>/examples/sql
> java -jar <KVHOME>/lib/sql.jar -helper-hosts <host>:<port> \
-store <storename> load \
-file <KVHOME>/examples/sql/SQLBasicExamples.cli

```

Running the SQL Shell

You can run SQL queries using the SQL shell. This is described in [Introduction to the SQL for Oracle NoSQL Database Shell](#). But, briefly, to run the queries shown in this document, start the shell:

```

java -jar KVHOME/lib/sql.jar
-helper-hosts node01:5000 -store kvstore
sql->

```

Note:

This book shows examples which are displayed in COLUMN mode. Be aware that the default output type is JSON. Use the `mode` command to switch between COLUMN and JSON (or JSON pretty) output.

Choosing column data

You can choose columns from a table. To do so, list the names of the desired table columns after `SELECT` in the statement, before noting the table after the `FROM` clause.

The `FROM` clause can name only one table. To retrieve data from a child table, use dot notation, such as `parent.child`.

To choose all table columns, use the asterisk (*) wildcard character as follows:

```
sql-> SELECT * FROM Users;
```

The `SELECT` statement displays these results:

```

+----+-----+-----+-----+-----+
| id | firstname | lastname | age | income |
+----+-----+-----+-----+-----+
| 3 | John      | Morgan   | 38  | NULL   |
| 4 | Peter     | Smith    | 38  | 80000  |
| 2 | John      | Anderson | 35  | 100000 |
| 5 | Dana      | Scully   | 47  | 400000 |

```

```
| 1 | David | Morrison | 25 | 100000 |
+-----+-----+-----+-----+
```

5 rows returned

To choose specific column(s) from the table Users, include the column names as a comma-separated list in the SELECT statement:

```
sql-> SELECT firstname, lastname, age FROM Users;
```

```
+-----+-----+-----+
| firstname | lastname | age |
+-----+-----+-----+
| John      | Morgan   | 38  |
| David     | Morrison | 25  |
| Dana     | Scully   | 47  |
| Peter    | Smith    | 38  |
| John     | Anderson | 35  |
+-----+-----+-----+
```

5 rows returned

Substituting column names for 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 returned data returned. In the next example, the query substitutes Surname for the actual column name lastname, by using the actual-name AS substitute-name clause, in the SELECT statement.

```
sql-> SELECT lastname AS Surname FROM Users;
```

```
+-----+
| Surname |
+-----+
| Scully  |
| Smith   |
| Morgan  |
| Anderson|
| Morrison|
+-----+
```

5 rows returned

Computing values for new columns

The SELECT statement can contain computational expressions based on the values of existing columns. For example, in the next statement, you select the values of one column, income, divide each value by 12, and display the output in another column. The SELECT statement can use almost any type of expression. If more than one value is returned, the items are inserted into an array.

This SELECT statement uses the yearly income values divided by 12 to calculate the corresponding values for monthsalary:

```
sql-> SELECT id, lastname, income, income/12
```

```
AS monthsalary FROM users;
```

```
+-----+-----+-----+-----+
| id | lastname | income | monthsalary |
+-----+-----+-----+-----+
```

2	Anderson	100000	8333
1	Morrison	100000	8333
5	Scully	400000	33333
4	Smith	80000	6666
3	Morgan	NULL	NULL

5 rows returned

This **SELECT** statement performs an addition operation that adds a bonus of 5000 to income to return `salarywithbonus`:

```
sql-> SELECT id, lastname, income, income+5000
AS salarywithbonus FROM users;
```

id	lastname	income	salarywithbonus
4	Smith	80000	85000
1	Morrison	100000	105000
5	Scully	400000	405000
3	Morgan	NULL	NULL
2	Anderson	100000	105000

5 rows returned

Identifying tables and their columns

The **FROM** clause can contain one table only (that is, joins are not supported). The table is specified by its name, which may be followed by an optional alias. The table can be referenced in the other clauses either by its name or its alias. As we will see later, sometimes the use of the table name or alias is mandatory. However, for table columns, the use of the table name or alias is optional. For example, here are three ways to write the same query:

```
sql-> SELECT Users.lastname, age FROM Users;
```

lastname	age
Scully	47
Smith	38
Morgan	38
Anderson	35
Morrison	25

5 rows returned

To identify the table `Users` with the alias `u`:

```
sql-> SELECT lastname, u.age FROM Users u ;
```

The keyword **AS** can optionally be used before an alias. For example, to identify the table `Users` with the alias `People`:

```
sql-> SELECT People.lastname, People.age FROM Users AS People;
```

Filtering results

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 comparison operators are also supported: =, !=, >, >=, <, and <= .

This query filters results to return only users whose first name is John:

```
sql-> SELECT id, firstname, lastname FROM Users WHERE firstname = "John";
+-----+-----+-----+
| id | firstname | lastname |
+-----+-----+-----+
| 3 | John      | Morgan   |
| 2 | John      | Anderson |
+-----+-----+-----+
```

2 rows returned

To return users whose calculated monthllysalary is greater than 6000:

```
sql-> SELECT id, lastname, income, income/12 AS monthllysalary
FROM Users WHERE income/12 > 6000;
+-----+-----+-----+-----+
| id | lastname | income | monthllysalary |
+-----+-----+-----+-----+
| 5 | Scully   | 400000 | 33333          |
| 4 | Smith    | 80000  | 6666           |
| 2 | Anderson | 100000 | 8333           |
| 1 | Morrison | 100000 | 8333           |
+-----+-----+-----+-----+
```

5 rows returned

To return users whose age is between 30 and 40 or whose income is greater than 100,000:

```
sql-> SELECT lastname, age, income FROM Users
WHERE age >= 30 and age <= 40 or income > 100000;
+-----+-----+-----+
| lastname | age | income |
+-----+-----+-----+
| Smith    | 38 | 80000  |
| Morgan   | 38 | NULL   |
| Anderson | 35 | 100000 |
| Scully   | 47 | 400000 |
+-----+-----+-----+
```

4 rows returned

You can use parenthesized expressions to alter the default precedence among operators. For example:

To return the users whose age is greater than 40 and either their age is less than 30 or their income is greater or equal than 100,000:

```
sql-> SELECT id, lastName FROM Users WHERE
(income >= 100000 or age < 30) and age > 40;
+-----+-----+
```

```

| id | lastName |
+----+-----+
| 5 | Scully   |
+----+-----+

```

1 row returned

You can use the `IS NULL` condition to return results where a field column value is set to `SQL NULL` (`SQL NULL` is used when a non-JSON field is set to null):

```

sql-> SELECT id, lastname from Users WHERE income IS NULL;
+----+-----+
| id | lastname |
+----+-----+
| 3 | Morgan   |
+----+-----+

```

1 row returned

You can use the `IS NOT NULL` condition to return column values that contain non-null data:

```

sql-> SELECT id, lastname from Users WHERE income IS NOT NULL;
+----+-----+
| id | lastname |
+----+-----+
| 4 | Smith    |
| 1 | Morrison |
| 5 | Scully   |
| 2 | Anderson |
+----+-----+

```

4 rows returned

Ordering Results

Use the `ORDER BY` clause to order the results by a primary key column or a non-primary key column.

Note:

You can use `ORDER BY` only if you are selecting by the table's primary key, or if there is an index that sorts the table's rows in the desired order.

To order by using a primary key column (`id`), specify the sort column in the `ORDER BY` clause:

```

sql-> SELECT id, lastname FROM Users ORDER BY id;
+----+-----+
| id | lastname |
+----+-----+
| 1 | Morrison |
| 2 | Anderson |
| 3 | Morgan   |
| 4 | Smith    |
| 5 | Scully   |
+----+-----+

```



```
+----+-----+
5 rows returned
```

To order by a non-primary key column, first create an index on the column of interest. For example, to use column lastname for ordering, create an index on that column, before using it in your ORDER BY clause:

```
sql-> CREATE INDEX idx1 on Users(lastname);
Statement completed successfully
sql-> SELECT id, lastname FROM Users ORDER BY lastname;
+----+-----+
| id | lastname |
+----+-----+
|  2 | Anderson |
|  3 | Morgan   |
|  1 | Morrison |
|  5 | Scully   |
|  4 | Smith    |
+----+-----+
```

```
5 rows returned
```

Using this example data, you can order by more than one column if you create an index on the columns. (If our table had used more than one column for its primary key, then you can order by multiple columns using the primary keys.) For example, to order users by age and income.

```
sql-> CREATE INDEX idx2 on Users(age, income);
Statement completed successfully
sql-> SELECT id, lastname, age, income FROM Users ORDER BY age, income;
+----+-----+-----+-----+
| id | lastname | age | income |
+----+-----+-----+-----+
|  1 | Morrison | 25 | 100000 |
|  2 | Anderson | 35 | 100000 |
|  4 | Smith    | 38 |  80000 |
|  3 | Morgan   | 38 |    NULL |
|  5 | Scully   | 47 | 400000 |
+----+-----+-----+-----+
```

```
5 rows returned
```

Creating a single index from two columns in the order you use them (age, income in this example), has some limits. The first column name (age) becomes the main sort item for the new index. You can use idx2 index to order by age only, but neither by income only, nor by income first and age second.

```
sql-> SELECT id, lastname, age from Users ORDER BY age;
+----+-----+-----+
| id | lastname | age |
+----+-----+-----+
|  1 | Morrison | 25 |
|  2 | Anderson | 35 |
|  4 | Smith    | 38 |
|  3 | Morgan   | 38 |
|  5 | Scully   | 47 |
+----+-----+-----+
```

```
5 rows returned
```

To learn more about indexes see [Working With Indexes](#).

By default, sorting is performed in ascending order. To sort in descending order use the DESC keyword in the ORDER BY clause:

```
sql-> SELECT id, lastname FROM Users ORDER BY id DESC;
+-----+-----+
| id | lastname |
+-----+-----+
| 5 | Scully   |
| 4 | Smith    |
| 3 | Morgan   |
| 2 | Anderson |
| 1 | Morrison |
+-----+-----+
```

5 rows returned

Limiting and Offsetting Results

Use the LIMIT clause to limit the number of results returned from a SELECT statement. For example, if there are 1000 rows in the Users table, limit the number of rows to return by specifying a LIMIT value. For example, this statement returns the first four ID rows from the table:

```
sql-> SELECT * from Users ORDER BY id LIMIT 4;
+-----+-----+-----+-----+-----+
| id | firstname | lastname | age | income |
+-----+-----+-----+-----+-----+
| 1 | David     | Morrison | 25 | 100000 |
| 2 | John      | Anderson | 35 | 100000 |
| 3 | John      | Morgan   | 38 | NULL   |
| 4 | Peter     | Smith    | 38 | 80000  |
+-----+-----+-----+-----+-----+
```

4 rows returned

To return only results 3 and 4 from the 10000 rows use the LIMIT clause to indicate 2 values, and the OFFSET clause to specify where the offset begins (after the first two rows). For example:

```
sql-> SELECT * from Users ORDER BY id LIMIT 2 OFFSET 2;
+-----+-----+-----+-----+-----+
| id | firstname | lastname | age | income |
+-----+-----+-----+-----+-----+
| 3 | John      | Morgan   | 38 | NULL   |
| 4 | Peter     | Smith    | 38 | 80000  |
+-----+-----+-----+-----+-----+
```

2 rows returned

Note:

We recommend using LIMIT and OFFSET with an ORDER BY clause. Otherwise, the results are returned in a random order, producing unpredictable results.

Using External Variables

Using external variables lets a query to be written and compiled once, and then run multiple times with different values for the external variables. Binding the external variables to specific values is done through APIs (see the *Getting Started with the Table API* manual), which you use before executing the query.

You must declare external variables in your SQL query before referencing them in the `SELECT` statement. For example:

```
DECLARE $age integer;
SELECT firstname, lastname, age
FROM Users
WHERE age > $age;
```

If the variable `$age` is set to value 39, the result of the above query is:

```
+-----+-----+-----+
| firstname | lastname | age |
+-----+-----+-----+
| Dana      | Scully   | 47  |
+-----+-----+-----+
```

3

Working with complex data

In this chapter, we present query examples that use complex data types (arrays, maps, records). To follow along with the examples, get the `Examples` download from [here](#) and run the `SQLAdvancedExamples` script found in the `sql` folder. This script creates the table and imports the data used.

SQLAdvancedExamples Script

The `SQLAdvancedExamples` script creates the following table:

```
CREATE TABLE Persons (  
  id integer,  
  firstname string,  
  lastname string,  
  age integer,  
  income integer,  
  lastLogin timestamp(4),  
  address record(street string,  
                 city string,  
                 state string,  
                 phones array(record(type enum(work, home),  
                                   areacode integer,  
                                   number integer)  
                               )  
                ),  
  connections array(integer),  
  expenses map(integer),  
  primary key (id)  
);
```

The script also imports the following table rows:

```
{  
  "id":1,  
  "firstname":"David",  
  "lastname":"Morrison",  
  "age":25,  
  "income":100000,  
  "lastLogin" : "2016-10-29T18:43:59.8319",  
  "address":{"street":"150 Route 2",  
            "city":"Antioch",  
            "state":"TN",  
            "zipcode" : 37013,  
            "phones":[{"type":"home", "areacode":423,  
                      "number":8634379}]  
            },  
  "connections":[2, 3],  
  "expenses":{"food":1000, "gas":180}  
}
```

```
{
  "id":2,
  "firstname":"John",
  "lastname":"Anderson",
  "age":35,
  "income":100000,
  "lastLogin" : "2016-11-28T13:01:11.2088",
  "address":{"street":"187 Hill Street",
    "city":"Beloit",
    "state":"WI",
    "zipcode" : 53511,
    "phones":[{"type":"home", "areacode":339,
      "number":1684972}]
  },
  "connections":[1, 3],
  "expenses":{"books":100, "food":1700, "travel":2100}
}

{
  "id":3,
  "firstname":"John",
  "lastname":"Morgan",
  "age":38,
  "income":100000000,
  "lastLogin" : "2016-11-29T08:21:35.4971",
  "address":{"street":"187 Aspen Drive",
    "city":"Middleburg",
    "state":"FL",
    "phones":[{"type":"work", "areacode":305,
      "number":1234079},
      {"type":"home", "areacode":305,
      "number":2066401}
    ]
  },
  "connections":[1, 4, 2],
  "expenses":{"food":2000, "travel":700, "gas":10}
}

{
  "id":4,
  "firstname":"Peter",
  "lastname":"Smith",
  "age":38,
  "income":80000,
  "lastLogin" : "2016-10-19T09:18:05.5555",
  "address":{"street":"364 Mulberry Street",
    "city":"Leominster",
    "state":"MA",
    "phones":[{"type":"work", "areacode":339,
      "number":4120211},
      {"type":"work", "areacode":339,
      "number":8694021},
      {"type":"home", "areacode":339,
      "number":1205678},
      {"type":"home", "areacode":305,
      "number":8064321}
    ]
  },
  "connections":[3, 5, 1, 2],
  "expenses":{"food":6000, "books":240, "clothes":2000, "shoes":1200}
}
```

```

{
  "id":5,
  "firstname":"Dana",
  "lastname":"Scully",
  "age":47,
  "income":400000,
  "lastLogin" : "2016-11-08T09:16:46.3929",
  "address":{"street":"427 Linden Avenue",
    "city":"Monroe Township",
    "state":"NJ",
    "phones":[{"type":"work", "areacode":201,
      "number":3213267},
      {"type":"work", "areacode":201,
      "number":8765421},
      {"type":"home", "areacode":339,
      "number":3414578}
    ]
  },
  "connections":[2, 4, 1, 3],
  "expenses":{"food":900, "shoes":1000, "clothes":1500}
}

```

You run the SQLAdvancedExamples script using the `load` command:

```

> cd <installdir>/examples/sql
> java -jar <KVHOME>/lib/sql.jar -helper-hosts <host>:<port> \
-store <storename> load \
-file <KVHOME>/examples/sql/SQLAdvancedExamples.cli

```

Note:

The Persons table schema models people that can be connected to other people in the table. All connections are stored in the "connections" column, which consists of an array of integers. Each integer is an ID of a person with whom the subject is connected. The entries in the "connections" array are sorted in descending order, indicating the strength of the connection. For example, looking at the record for person 3, we see that John Morgan has these connections: [1, 4, 2]. The order of the array elements specifies that John is most strongly connected with person 1, less connected with person 4, and least connected with person 2.

Records in the Persons table also include an "expenses" column, declared as an integer map. For each person, the map stores key-value pairs of string item types and integers representing money spent on the item. For example, one record has these expenses: {"food":900, "shoes":1000, "clothes":1500}, other records have different items. One benefit of modelling expenses as a map type is to facilitate the categories being different for each person. Later, we may want to add or delete categories dynamically, without changing the table schema, which maps readily support. An item to note about this map is that it is an integer map always contains key-value pairs, and keys are always strings.

Working with Timestamps

To specify a timestamp value in a query, provide it as a string, and cast it to a Timestamp data type. For example:

```
sql-> SELECT id, firstname, lastname FROM Persons WHERE
lastLogin = CAST("2016-10-19T09:18:05.5555" AS TIMESTAMP);
```

id	firstname	lastname
4	Peter	Smith

1 row returned

Timestamp queries often involve a range of time, which requires multiple casts:

```
sql-> SELECT id, firstname, lastname, lastLogin FROM Persons WHERE
lastLogin > CAST("2016-11-01" AS TIMESTAMP) AND
lastLogin < CAST("2016-11-30" AS TIMESTAMP);
```

id	firstname	lastname	lastLogin
3	John	Morgan	2016-11-29T08:21:35.4971
2	John	Anderson	2016-11-28T13:01:11.2088
5	Dana	Scully	2016-11-08T09:16:46.3929

3 rows returned

You can also use various Timestamp functions to return specific time and date values from the Timestamp data. (See [Timestamp Functions](#) for a full list of these functions). For example:

```
sql-> SELECT id, firstname, lastname,
year(lastLogin) AS Year,
month(lastLogin) AS Month,
day(lastLogin) AS Day,
hour(lastLogin) AS Hour,
minute(lastLogin) AS Minute
FROM Persons;
```

id	firstname	lastname	Year	Month	Day	Hour	Minute
3	John	Morgan	2016	11	29	8	21
2	John	Anderson	2016	11	28	13	1
4	Peter	Smith	2016	10	19	9	18
5	Dana	Scully	2016	11	8	9	16
1	David	Morrison	2016	10	29	18	43

Alternatively, use the EXTRACT function:

```
sql-> SELECT id, firstname, lastname,
EXTRACT(YEAR FROM lastLogin) AS Year,
EXTRACT(MONTH FROM lastLogin) AS Month,
EXTRACT(DAY FROM lastLogin) AS Day,
EXTRACT(HOUR FROM lastLogin) AS Hour,
EXTRACT(MINUTE FROM lastLogin) AS Minute
```

```
FROM Persons;
```

id	firstname	lastname	Year	Month	Day	Hour	Minute
3	John	Morgan	2016	11	29	8	21
4	Peter	Smith	2016	10	19	9	18
1	David	Morrison	2016	10	29	18	43
2	John	Anderson	2016	11	28	13	1
5	Dana	Scully	2016	11	8	9	16

```
5 rows returned
sql->
```

Working With Arrays

You can use slice or filter steps to select elements out of an array. We start with some examples using slice steps.

To select and display the second connection of each person, we use this query:

```
sql-> SELECT lastname, connections[1]
AS connection FROM Persons;
```

lastname	connection
Scully	2
Smith	4
Morgan	2
Anderson	2
Morrison	2

```
5 rows returned
```

In the example, the slice step [1] is applied to the connections array. Since array elements start with 0, 1 selects the second connection value.

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.

For example, the following query returns the lastname and the first 3 connections of person 5 as strongconnections:

```
sql-> SELECT lastname, [connections[0:2]]
AS strongconnections FROM Persons WHERE id = 5;
```

lastname	strongconnections
Scully	2 4 1

```
1 row returned
```

In the above query for Person 5, the path expression `connections[0:2]` returns the person's first 3 connections. Here, the range is [0:2], so 0 is the low expression and 2 is the high. The path expression returns its result as a list of 3 items. The list is

converted to an array (a single item) by enclosing the path expression in an array-constructor expression (`[]`). The array constructor creates a new array containing the three connections. Notice that although the query shell displays the elements of this constructed array vertically, the number of rows returned by this query is 1.

Use of the array constructor in the select clause is optional. If no array constructor is used, an array will still be constructed, but only if the select-clause expression does indeed return more than one item. If exactly one item is returned, the result will contain just that one item. If the expression returns nothing (an empty result), NULL is used as the result. This behavior is illustrated in the next example, which we will run with and without an array constructor.

As mentioned above, you can omit the low or high expression when specifying the range for a slice step. For example the following query specifies a range of `[3:]` which returns all connections after the third one. Notice that for persons having only 3 connections or less, an empty array is constructed and returned due to the use of the array constructor.

To fully illustrate this behavior, we display this output in mode `JSON` because the `COLUMN` mode does not differentiate between a single item and an array containing a single item.

```
sql-> mode JSON
Query output mode is JSON
sql-> SELECT id, [connections[3:]] AS weakConnections FROM Persons;
{"id":3,"weakConnections":[]}
{"id":4,"weakConnections":[2]}
{"id":2,"weakConnections":[]}
{"id":5,"weakConnections":[3]}
{"id":1,"weakConnections":[]}
```

5 rows returned

Now we run the same query, but without the array constructor. Notice how single items are not contained in an array, and for rows with no match, NULL is returned instead of an empty array.

```
sql-> SELECT id, connections[3:] AS weakConnections FROM Persons;
{"id":2,"weakConnections":null}
{"id":3,"weakConnections":null}
{"id":4,"weakConnections":2}
{"id":5,"weakConnections":3}
{"id":1,"weakConnections":null}
```

5 rows returned

```
sql-> mode COLUMN
Query output mode is COLUMN
sql->
```

As a last example of slice steps, the following query returns the last 3 connections of each person. In this query, the slice step is `[size($)-3:]`. In this expression, the `$` is an implicitly declared variable that references the array that the slice step is applied to. In this example, `$` references the connections array. The `size()` built-in function returns the size (number of elements) of the input array. So, in this example, `size($)` is the size of the current connections array. Finally, `size($)-3` computes the third position from the end of the current connections array.

```
sql-> SELECT id, [connections[size($)-3:]]
AS weakConnections FROM Persons;
+-----+-----+
```

id	weakConnections
5	4 1 3
4	5 1 2
3	1 4 2
2	1 3
1	2 3

5 rows returned

We now turn our attention to filter steps on arrays. Like slice steps, filter steps also use the square brackets ([]) syntax. However, what goes inside the [] is different. With filter steps there is either nothing inside the [] or a single expression that acts as a condition (returns a boolean result). In the former case, all the elements of the array are selected (the array is "unnested"). In the latter case, the condition is applied to each element in turn, and if the result is true, the element is selected, otherwise it is skipped. For example:

The following query returns the id and connections of persons who are connected to person 4:

```
sql-> SELECT id, connections
FROM Persons p WHERE p.connections[] =any 4;
```

id	connections
3	1 4 2
5	2 4 1 3

2 rows returned

In the above query, the expression `p.connections[]` returns all the connections of a person. Then, the `=any` operator returns true if this sequence of connections contains the number 4. Sequence operators are described in [Sequence Comparison Operators](#).

The following query returns the id and connections of persons who are connected with any person having an id greater than 4:

```
sql-> SELECT id, connections FROM Persons p
WHERE p.connections[] >any 4;
-----+-----
```

```

| id | connections |
+---+-----+
|  4 | 3           |
|    | 5           |
|    | 1           |
|    | 2           |
+---+-----+

```

1 row returned

The following query returns, for each person, the person's last name and the phone numbers with area code 339:

```

sql-> SELECT lastname,
[ p.address.phones[$element.areacode = 339].number ]
AS phoneNumbers FROM Persons p;

```

```

+-----+-----+
| lastname | phoneNumbers |
+-----+-----+
| Scully   | 3414578      |
+-----+-----+
| Smith    | 4120211      |
|          | 8694021      |
|          | 1205678      |
+-----+-----+
| Morgan   |               |
+-----+-----+
| Anderson | 1684972      |
+-----+-----+
| Morrison |               |
+-----+-----+

```

5 rows returned

In the above query, the filter step `[$element.areacode = 339]` is applied to the `phones` array of each person. The filter step evaluates the condition `$element.areacode = 339` on each element of the array. This condition expression uses the implicitly declared variable `$element`, which references the current element of the array. An empty array is returned for persons that do not have any phone number in the 339 area code. If we wanted to filter out such persons from the result, we would write the following query:

```

sql-> SELECT lastname,
[ p.address.phones[$element.areacode = 339].number ]
AS phoneNumbers FROM Persons p WHERE p.address.phones.areacode =any 339;

```

```

+-----+-----+
| lastname | phoneNumbers |
+-----+-----+
| Scully   | 3414578      |
+-----+-----+
| Smith    | 4120211      |
|          | 8694021      |
|          | 1205678      |
+-----+-----+
| Anderson | 1684972      |
+-----+-----+

```

3 rows returned

The previous query contains the path expression `p.address.phones.areacode`. In that expression, the field step `.areacode` is applied to an array field (`phones`). In this case,

the field step is applied to each element of the array in turn. In fact, the path expression is equivalent to `p.address.phones[].areacode`.

In addition to the implicitly-declared `$` and `$element` variables, the condition inside a filter step can also use the `$pos` variable (also implicitly declared). `$pos` references the position within the array of the current element (the element on which the condition is applied). For example, the following query selects the "interesting" connections of each person, where a connection is considered interesting if it is among the 3 strongest connections and connects to a person with an id greater or equal to 4.

```
sql-> SELECT id, [p.connections[$element >= 4 and $pos < 3]]
AS interestingConnections FROM Persons p;
+-----+-----+-----+
| id | interestingConnections |
+-----+-----+-----+
| 5 | 4 |
+-----+-----+-----+
| 4 | 5 |
+-----+-----+-----+
| 3 | 4 |
+-----+-----+-----+
| 2 | |
+-----+-----+-----+
| 1 | |
+-----+-----+-----+
```

5 rows returned

Finally, two arrays can be compared with each other using the usual comparison operators (`=`, `!=`, `>`, `>=`, `<`, and `<=`). For example the following query constructs the array `[1,3]` and selects persons whose connections array is equal to `[1,3]`.

```
sql-> SELECT lastname FROM Persons p
WHERE p.connections = [1,3];
+-----+
| lastname |
+-----+
| Anderson |
+-----+
```

1 row returned

Working with Records

You can use a field step to select the value of a field from a record. For example, to return the id, last name, and city of persons who reside in Florida:

```
sql-> SELECT id, lastname, p.address.city
FROM Persons p WHERE p.address.state = "FL";
+-----+-----+-----+
| id | lastname | city |
+-----+-----+-----+
| 3 | Morgan | Middleburg |
+-----+-----+-----+
```

1 row returned

In the above query, the path expression (see [Path Expressions](#)) `p.address.state` consists of 2 field steps: `.address` selects the address field of the current row (rows can

be viewed as records, whose fields are the row columns), and `.state` selects the state field of the current address.

The example record contains an array of phone numbers. You can form queries against that array using a combination of path steps and sequence comparison operators (see [Sequence Comparison Operators](#)). For example, to return the last name of persons who have a phone number with area code 423:

```
sql-> SELECT lastname FROM Persons
p WHERE p.address.phones.areacode =any 423;
+-----+
| lastname |
+-----+
| Morrison |
+-----+
```

1 row returned

In the above query, the path expression `p.address.phones.areacode` returns all the area codes of a person. Then, the `=any` operator returns true if this sequence of area codes contains the number 423. Notice also that the field step `.areacode` is applied to an array field (`phones`). This is allowed if the array contains records or maps. In this case, the field step is applied to each element of the array in turn.

The following example returns all the persons who had three connections. Notice the use of `[]` after `connections`: it is an array filter step, which returns all the elements of the `connections` array as a sequence (it is unnesting the array).

```
sql-> SELECT id, firstName, lastName, connections from Persons where
connections[] =any 3 ORDER BY id;
+-----+-----+-----+-----+
| id | firstName | lastName | connections |
+-----+-----+-----+-----+
| 1 | David     | Morrison | 2            |
|   |           |         | 3            |
+-----+-----+-----+-----+
| 2 | John     | Anderson | 1            |
|   |           |         | 3            |
+-----+-----+-----+-----+
| 4 | Peter    | Smith   | 3            |
|   |           |         | 5            |
|   |           |         | 1            |
|   |           |         | 2            |
+-----+-----+-----+-----+
| 5 | Dana     | Scully  | 2            |
|   |           |         | 4            |
|   |           |         | 1            |
|   |           |         | 3            |
+-----+-----+-----+-----+
```

4 rows returned

This query can use `ORDER BY` to sort the results because the sort is being performed on the table's primary key. The next section shows sorting on non-primary key fields through the use of indexes.

See [Working With Arrays](#) for more examples of querying against data contained in arrays.

Using ORDER BY to Sort Results

To sort the results from a SELECT statement using a field that is not the table's primary key, you must first create an index for the column of choice. For example, for the next table, to query based on a Timestamp and sort the results in descending order by the timestamp, create an index:

```
sql-> SELECT id, firstname, lastname, lastLogin FROM Persons;
```

id	firstname	lastname	lastLogin
3	John	Morgan	2016-11-29T08:21:35.4971
4	Peter	Smith	2016-10-19T09:18:05.5555
2	John	Anderson	2016-11-28T13:01:11.2088
5	Dana	Scully	2016-11-08T09:16:46.3929
1	David	Morrison	2016-10-29T18:43:59.8319

5 rows returned

```
sql-> CREATE INDEX tsidx1 on Persons (lastLogin);
```

Statement completed successfully

```
sql-> SELECT id, firstname, lastname, lastLogin
FROM Persons ORDER BY lastLogin DESC;
```

id	firstname	lastname	lastLogin
3	John	Morgan	2016-11-29T08:21:35.4971
2	John	Anderson	2016-11-28T13:01:11.2088
5	Dana	Scully	2016-11-08T09:16:46.3929
1	David	Morrison	2016-10-29T18:43:59.8319
4	Peter	Smith	2016-10-19T09:18:05.5555

5 rows returned

SQL for Oracle NoSQL Database can also sort query results by the values of nested records. To do so, create an index of the nested field (or fields). For example, you can create an index of address.state from the Persons table, and then order by state:

```
sql-> CREATE INDEX idx1 on Persons (address.state);
```

Statement completed successfully

```
sql-> SELECT id, $p.address.state FROM
Persons $p ORDER BY $p.address.state;
```

id	state
3	FL
4	MA
5	NJ
1	TN
2	WI

5 rows returned

To learn more about indexes see [Working With Indexes](#).

Working With Maps

The path steps applicable to maps are field and filter steps. Slice steps do not make sense for maps, because maps are unordered, and as a result, their entries do not have any fixed positions.

You can use a field step to select the value of a field from a map. For example, to return the lastname and the food expenses of all persons:

```
sql-> SELECT lastname, p.expenses.food
FROM Persons p;
```

```
+-----+-----+
| lastname | food |
+-----+-----+
| Morgan   | 2000 |
| Morrison | 1000 |
| Scully   | 900  |
| Smith    | 6000 |
| Anderson | 1700 |
+-----+-----+
```

5 rows returned

In the above query, the path expression `p.expenses.food` consists of 2 field steps: `.expenses` selects the expenses field of the current row and `.food` selects the value of the food field/entry from the current expenses map.

To return the lastname and amount spent on travel for each person who spent less than \$3000 on food:

```
sql-> SELECT lastname, p.expenses.travel
FROM Persons p WHERE p.expenses.food < 3000;
```

```
+-----+-----+
| lastname | travel |
+-----+-----+
| Scully   | NULL  |
| Morgan   | 700   |
| Anderson | 2100  |
| Morrison | NULL  |
+-----+-----+
```

4 rows returned

Notice that NULL is returned for persons who did not have any travel expenses.

Filter steps are performed using either the `.values()` or `.keys()` path steps. To select values of map entries, use `.values(<cond>)`. To select keys of map entries, use `.keys(<cond>)`. If no condition is used in these steps, all the values or keys of the input map are selected. If the steps do contain a condition expression, the condition is evaluated for each entry, and the value or key of the entry is selected/skipped if the result is true/false.

The implicitly-declared variables `$key` and `$value` can be used inside a map filter condition. `$key` references the key of the current entry and `$value` references the associated value. Notice that, contrary to arrays, the `$pos` variable can not be used inside map filters (because map entries do not have fixed positions).

To show, for each user, their id and the expense categories where they spent more than \$1000:

```
sql-> SELECT id, p.expenses.keys($value > 1000) as Expenses
from Persons p;
```

id	Expenses
4	clothes food shoes
3	food
2	food travel
5	clothes
1	NULL

To return the id and the expense categories in which the user spent more than they spent on clothes, use the following filter step expression. In this query, the context-item variable (\$) appearing in the filter step expression [\$value > \$.clothes] refers to the expenses map as a whole.

```
sql-> SELECT id, p.expenses.keys($value > $.clothes) FROM Persons p;
```

id	Column_2
3	NULL
2	NULL
5	NULL
1	NULL
4	food

To return the id and expenses data of any person who spent more on any category than what they spent on food:

```
sql-> SELECT id, p.expenses
FROM Persons p
WHERE p.expenses.values() >any p.expenses.food;
```

id	expenses
5	clothes 1500 food 900 shoes 1000
2	books 100 food 1700 travel 2100

2 rows returned

To return the id of all persons who consumed more than \$2000 in any category other than food:

```
sql-> SELECT id FROM Persons p
WHERE p.expenses.values($key != "food") >any 2000;
+----+
| id |
+----+
| 2 |
+----+
```

1 row returned

Using the size() Function

The size function can be used to return the size (number of fields/entries) of a complex item (record, array, or map). For example:

To return the id and the number of phones that each person has:

```
sql-> SELECT id, size(p.address.phones)
AS registeredphones FROM Persons p;
+-----+-----+
| id | registeredphones |
+-----+-----+
| 5 | 3 |
| 3 | 2 |
| 4 | 4 |
| 2 | 1 |
| 1 | 1 |
+-----+-----+
```

5 rows returned

To return the id and the number of expenses categories for each person: has:

```
sql-> SELECT id, size(p.expenses) AS
categories FROM Persons p;
+-----+-----+
| id | categories |
+-----+-----+
| 4 | 4 |
| 3 | 3 |
| 2 | 3 |
| 1 | 2 |
| 5 | 3 |
+-----+-----+
```

5 rows returned

To return for each person their id and the number of expenses categories for which the expenses were more than 2000:

```
sql-> SELECT id, size([p.expenses.values($value > 2000)]) AS
expensiveCategories FROM Persons p;
+-----+-----+
| id | expensiveCategories |
+-----+-----+
| 3 | 0 |
| 2 | 1 |
```

```
| 5 |           0 |  
| 1 |           0 |  
| 4 |           1 |  
+---+-----+  
|
```

5 rows returned

4

Working with JSON

This chapter provides examples on working with JSON data. If you want to follow along with the examples, get the `Examples` download from here and run the `SQLJSONExamples` script found in the `sql` folder. This creates the table and imports the data used.

JSON data is written to JSON data columns by providing a JSON object. This object can contain any valid JSON data. The input data is parsed and stored internally as Oracle NoSQL Database datatypes:

- When numbers are encountered, they are converted to integer, long, or double items, depending on the actual value of the number (float items are not used for JSON).
- Strings in the input text are mapped to string items.
- Boolean values are mapped to boolean items.
- JSON nulls are mapped to JSON null items.
- When an array is encountered in the input text, an array item is created whose type is `Array(JSON)`. This is done unconditionally, no matter what the actual contents of the array might be.
- When a JSON object is encountered in the input text, a map item is created whose type is `Map(JSON)`, unconditionally.

Note:

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

The remainder of this chapter provides an overview to querying JSON data.

SQLJSONExamples Script

The `SQLJSONExample` is available to illustrate JSON usage. This script creates the following table:

```
create table if not exists JSONPersons (  
  id integer,  
  person JSON,  
  primary key (id)  
);
```

The script imports the following table rows. Notice that the content for the `person` column, which is of type `JSON` contains a JSON object. That object contains a series of fields which represent our person. We have deliberately included inconsistent

information in this example so as to illustrate how to handle various queries when working with JSON data.

```
{
  "id":1,
  "person" : {
    "firstname":"David",
    "lastname":"Morrison",
    "age":25,
    "income":100000,
    "lastLogin" : "2016-10-29T18:43:59.8319",
    "address":{"street":"150 Route 2",
      "city":"Antioch",
      "state":"TN",
      "zipcode" : 37013,
      "phones":[{"type":"home", "areacode":423,
        "number":8634379}]
    },
    "connections":[2, 3],
    "expenses":{"food":1000, "gas":180}
  }
}

{
  "id":2,
  "person" : {
    "firstname":"John",
    "lastname":"Anderson",
    "age":35,
    "income":100000,
    "lastLogin" : "2016-11-28T13:01:11.2088",
    "address":{"street":"187 Hill Street",
      "city":"Beloit",
      "state":"WI",
      "zipcode" : 53511,
      "phones":[{"type":"home", "areacode":339,
        "number":1684972}]
    },
    "connections":[1, 3],
    "expenses":{"books":100, "food":1700, "travel":2100}
  }
}

{
  "id":3,
  "person" : {
    "firstname":"John",
    "lastname":"Morgan",
    "age":38,
    "income":100000000,
    "lastLogin" : "2016-11-29T08:21:35.4971",
    "address":{"street":"187 Aspen Drive",
      "city":"Middleburg",
      "state":"FL",
      "phones":[{"type":"work", "areacode":305,
        "number":1234079},
        {"type":"home", "areacode":305,
        "number":2066401}
      ]
    },
    "connections":[1, 4, 2],
```

```
        "expenses":{"food":2000, "travel":700, "gas":10}
    }
}
{
  "id":4,
  "person": {
    "firstname":"Peter",
    "lastname":"Smith",
    "age":38,
    "income":80000,
    "lastLogin" : "2016-10-19T09:18:05.5555",
    "address":{"street":"364 Mulberry Street",
      "city":"Leominster",
      "state":"MA",
      "phones":[{"type":"work", "areacode":339,
        "number":4120211},
        {"type":"work", "areacode":339,
        "number":8694021},
        {"type":"home", "areacode":339,
        "number":1205678},
        null,
        {"type":"home", "areacode":305,
        "number":8064321}
      ]
    },
    "connections":[3, 5, 1, 2],
    "expenses":{"food":6000, "books":240, "clothes":2000,
      "shoes":1200}
  }
}
{
  "id":5,
  "person" : {
    "firstname":"Dana",
    "lastname":"Scully",
    "age":47,
    "income":400000,
    "lastLogin" : "2016-11-08T09:16:46.3929",
    "address":{"street":"427 Linden Avenue",
      "city":"Monroe Township",
      "state":"NJ",
      "phones":[{"type":"work", "areacode":201,
        "number":3213267},
        {"type":"work", "areacode":201,
        "number":8765421},
        {"type":"home", "areacode":339,
        "number":3414578}
      ]
    },
    "connections":[2, 4, 1, 3],
    "expenses":{"food":900, "shoes":1000, "clothes":1500}
  }
}
{
  "id":6,
  "person" : {
    "mynumber":5,
    "myarray":[1,2,3,4]
  }
}
```

```

    }
  }
  {
    "id":7,
    "person" : {
      "mynumber":"5",
      "myarray":["1","2","3","4"]
    }
  }
}

```

You run the SQLJSONExamples script using the [load](#) command:

```

> cd <installdir>/examples/sql
> java -jar <KVHOME>/lib/sql.jar -helper-hosts <host>:<port> \
-store <storename> load \
-file <KVHOME>/examples/sql/SQLJSONExamples.cli

```

Basic Queries

Because JSON is parsed and stored internally in native data formats with Oracle NoSQL Database, querying JSON data is no different than querying data in other column types. See [Simple SELECT Queries](#) and [Working with complex data](#) for introductory examples of how to form these queries.

In our JSONPersons example, all of the data for each person is contained in a column of type JSON called `person`. This data is presented as a JSON object, and mapped internally into a `Map(JSON)` type. You can query information in this column as you would query a `Map` of any other type. For example:

```
sql-> SELECT id, j.person.lastname, j.person.age FROM JSONPersons j;
```

id	lastname	age
3	Morgan	38
2	Anderson	35
5	Scully	47
1	Morrison	25
4	Smith	38
6	NULL	NULL
7	NULL	NULL

7 rows returned

The last two rows in returned from this query contain all NULLs. This is because those rows were populated using JSON objects that are different than the objects used to populate the rest of the table. This capability of JSON is both a strength and a weakness. As a plus, you can modify your schema easily. However, if you are not careful, you can end up with tables containing dissimilar data in both large and small ways.

Because the JSON object is stored as a map, you can use normal map step functions on the column. For example:

```
sql-> SELECT id, j.person.expenses.keys($value > 1000) as Expenses
from JSONPersons j;
```

id	Expenses
3	food
2	food travel
4	clothes food shoes
6	NULL
5	clothes
7	NULL
1	NULL

7 rows returned

Here, id 1 is NULL because that user had no expenses greater than \$1000, while id 6 and 7 are NULL because they have no `j.person.expenses` field.

Using WHERE EXISTS with JSON

As we saw in the previous section, different rows in the same table can have dissimilar information in them when a column type is JSON. To identify whether desired information exists for a given JSON column, use the `EXISTS` operator.

For example, some of the JSON persons have a zip code entered for their address, and others do not. Use this query to see all the users with a zipcode:

```
sql-> SELECT id, j.person.address AS Address FROM JSONPersons j
WHERE EXISTS j.person.address.zipcode;
```

id	Address
2	city Beloit phones areacode 339 number 1684972 type home state WI street 187 Hill Street zipcode 53511
1	city Antioch phones areacode 423 number 8634379 type home state TN

```

| | street      | 150 Route 2 |
| | zipcode    | 37013      |
+-----+

```

2 rows returned

When querying data for inconsistencies, it is often more useful to see all rows where information is missing by using `WHERE NOT EXISTS`:

```
sql-> SELECT * FROM JSONPersons j WHERE NOT EXISTS j.person.lastname;
```

```

+-----+
| id | person      |
+-----+
| 7  | myarray     |
|    |             | 1
|    |             | 2
|    |             | 3
|    |             | 4
|    | mynumber   | 5
+-----+
| 6  | myarray     |
|    |             | 1
|    |             | 2
|    |             | 3
|    |             | 4
|    | mynumber   | 5
+-----+

```

1 row returned

Seeking NULLS in Arrays

All arrays found in a JSON input stream are stored internally as `ARRAY(JSON)`. This means that it is possible for the array to have inconsistent types for its members.

In our example, the phones array for user id 4 contains a null element:

```
sql-> SELECT j.person.address.phones FROM JSONPersons j WHERE j.id=4;
```

```

+-----+
| phones |
+-----+
| areacode | 339 |
| number   | 4120211 |
| type     | work |
+-----+
| areacode | 339 |
| number   | 8694021 |
| type     | work |
+-----+
| areacode | 339 |
| number   | 1205678 |
| type     | home |
| null     |      |
+-----+
| areacode | 305 |
| number   | 8064321 |
| type     | home |
+-----+

```


A way to discover this in your table is to examine the phones array for null values:

```
sql-> SELECT id, j.person.address.phones FROM JSONPersons j
WHERE j.person.address.phones[] =any null;
```

id	phones
4	areacode 339 number 4120211 type work
	areacode 339 number 8694021 type work
	areacode 339 number 1205678 type home null
	areacode 305 number 8064321 type home

1 row returned

Notice the use of the array filter step ([]) in the previous query. This is needed to unpack the array into a sequence so that the =any comparison operator can be used with it.

Examining Data Types JSON Columns

The example data contains a couple of rows with unusual data:

```
{
  "id":6,
  "person" : {
    "mynumber":5,
    "myarray":[1,2,3,4]
  }
}

{
  "id":7,
  "person" : {
    "mynumber":"5",
    "myarray":["1","2","3","4"]
  }
}
```

You can locate them using the query:

```
sql-> SELECT * FROM JSONPersons j WHERE EXISTS j.person.mynumber;
```

id	person
6	myarray
	1
	2

		3	
		4	
	mynumber	5	
+-----+			
7	myarray		
		1	
		2	
		3	
		4	
	mynumber	5	
+-----+			

2 rows returned

However, notice that these two rows actually contain numbers stored as different types. ID 6 stores integers while ID 7 stores strings. You can select a row based on its type:

```
sql-> SELECT * FROM JSONPersons j
WHERE j.person.mynumber IS OF TYPE (integer);
```

id	person
6	myarray
	1
	2
	3
	4
	mynumber 5

Notice that if you use `IS NOT OF TYPE` then every row in the table is returned except id 6. This is because for all the other rows, `j.person.mynumber` evaluates to `jnull`, which is not an integer.

```
sql-> SELECT id FROM JSONPersons j
WHERE j.person.mynumber IS NOT OF TYPE (integer);
```

id
3
2
5
4
1
7

6 rows returned

To solve this problem, also check for the existence of `j.person.mynumber`:

```
sql-> SELECT id from JSONPersons j WHERE EXISTS j.person.mynumber
and j.person.mynumber IS NOT OF TYPE (integer);
```

id
7

1 row returned

You can also perform type checking based on the type of data contained in the array. Recall that our rows contain arrays with integers and arrays with strings. You can return the row with just the array of strings using:

```
sql-> SELECT id, j.person.myarray FROM JSONPersons j
WHERE j.person.myarray[] IS OF TYPE (string+);
```

id	myarray
7	1
	2
	3
	4

1 row returned

Here, we use the array filter step (`[]`) in the WHERE clause to unpack the array into a sequence. This allows is-of-type to iterate over the sequence, checking the type of each element. If every element in the sequence matches the identified type (`string`, in this case), then the is-of-type returns true.

Also notice that the query uses the `+` cardinality modifier. This means that is-of-type will return true only if the input sequence (`myarray[]`, in this case) contains ONE OR MORE elements that match the identified type (`string`). If we used `*`, then 0 or more elements would have to match the identified type in order for true to return. Because our table contains a mix of rows with different schema, the result is that every row except id 6 is returned:

```
sql-> SELECT id, j.person.myarray FROM JSONPersons j
WHERE j.person.myarray[] IS OF TYPE (string*);
```

id	myarray
3	NULL
5	NULL
1	NULL
7	1
	2
	3
	4
4	NULL
2	NULL

6 rows returned

Finally, if we do not provide a cardinality modifier at all, then is-of-type returns true if ONE AND ONLY one member of the input sequence matches the identified type. In this example, the result is that no rows are returned.

```
sql-> SELECT id, j.person.myarray FROM JSONPersons j
WHERE j.person.myarray[] IS OF TYPE (string);
```

0 row returned

Using Map Steps with JSON Data

On import, Oracle NoSQL Database stores JSON objects as MAP(JSON). This means you can use map filter steps with your JSON objects.

For example, if you want to visually examine the JSON fields in use by your rows:

```
sql-> SELECT id, j.person.keys() FROM JSONPersons j;
```

id	Column_2
4	address age connections expenses firstname income lastLogin lastname
6	myarray mynumber
3	address age connections expenses firstname income lastLogin lastname
5	address age connections expenses firstname income lastLogin lastname
1	address age connections expenses firstname income lastLogin lastname
7	myarray mynumber
2	address age

```

|      | connections |
|      | expenses   |
|      | firstname  |
|      | income     |
|      | lastLogin  |
|      | lastname   |
+-----+

```

7 rows returned

Casting Datatypes

You can cast one data type to another using the `cast` expression. See [Cast Expressions](#) for rules about supported data type casting.

In JSON, casting is particularly useful for timestamp information because JSON has no equivalent to the Oracle NoSQL Database Timestamp data type. Instead, the timestamp information is carried in a JSON object as a string. To work with it as a Timestamp, use `cast`.

In [Working with Timestamps](#) we showed how to work with the timestamp data type. In this case, what you do is no different except you must cast both sides of the expression. Also, because the left side of the expression is a sequence, you must specify a type quantifier (* in this case):

```

sql-> SELECT id,
          j.person.firstname, j.person.lastname, j.person.lastLogin
       FROM JSONPersons j
       WHERE CAST(j.person.lastLogin AS TIMESTAMP*) >
             CAST("2016-11-01" AS TIMESTAMP) AND
             CAST(j.person.lastLogin AS TIMESTAMP*) <
             CAST("2016-11-30" AS TIMESTAMP);

```

```

+-----+-----+-----+-----+
| id | firstname | lastname | lastLogin |
+-----+-----+-----+-----+
| 3 | John      | Morgan   | 2016-11-29T08:21:35.4971 |
+-----+-----+-----+-----+
| 2 | John      | Anderson | 2016-11-28T13:01:11.2088 |
+-----+-----+-----+-----+
| 5 | Dana      | Scully   | 2016-11-08T09:16:46.3929 |
+-----+-----+-----+-----+

```

3 rows returned

As another example, you can cast to an integer and then operate on that number:

```

sql-> SELECT id, j.person.mynumber,
          CAST(j.person.mynumber as integer) * 10 AS TenTimes
       FROM JSONPersons j WHERE EXISTS j.person.mynumber;

```

```

+-----+-----+-----+
| id | mynumber | TenTimes |
+-----+-----+-----+
| 7 | 5         | 50       |
+-----+-----+-----+
| 6 | 5         | 50       |
+-----+-----+-----+

```

If you want to operate on just the row that contains the number as a string, use `IS OF TYPE`:

```

sql-> SELECT id, j.person.mynumber,
          CAST(j.person.mynumber as integer) * 10 AS TenTimes
FROM JSONPersons j WHERE EXISTS j.person.mynumber
AND j.person.mynumber IS OF TYPE (string);

```

id	mynumber	TenTimes
7	5	50

Using Searched Case

A searched case expression can be helpful in identifying specific problems with the JSON data in your JSON columns. The example data we have been using in this chapter sometimes provides a `JSONPersons.address` field, and sometimes it does not. When an address is present, sometimes it provides a `zipcode`, and sometimes it does not. We can use a searched case expression to identify and describe the specific problem with each row.

```

sql-> SELECT id,
CASE
  WHEN NOT EXISTS j.person.address
  THEN j.person.keys()
  WHEN NOT EXISTS j.person.address.zipcode
  THEN "No Zipcode"
  ELSE j.person.address.zipcode
END
FROM JSONPersons j;

```

id	Column_2
4	No Zipcode
3	No Zipcode
5	No Zipcode
1	37013
7	myarray mynumber
6	myarray mynumber
2	53511

7 rows returned

We can improve the report by adding a third column that uses a second searched case expression:

```

sql-> SELECT id,
CASE
  WHEN NOT EXISTS j.person.address
  THEN "No Address"
  WHEN NOT EXISTS j.person.address.zipcode
  THEN "No Zipcode"
  ELSE j.person.address.zipcode

```

```

END,
CASE
  WHEN NOT EXISTS j.person.address
  THEN j.person.keys()
  ELSE j.person.address
END
FROM JSONPersons j;

```

id	Column_2	Column_3
3	No Zipcode	city Middleburg phones areacode 305 number 1234079 type work areacode 305 number 2066401 type home state FL street 187 Aspen Drive
2	53511	city Beloit phones areacode 339 number 1684972 type home state WI street 187 Hill Street zipcode 53511
5	No Zipcode	city Monroe Township phones areacode 201 number 3213267 type work areacode 201 number 8765421 type work areacode 339 number 3414578 type home state NJ street 427 Linden Avenue
1	37013	city Antioch phones areacode 423 number 8634379 type home state TN street 150 Route 2 zipcode 37013
7	No Address	myarray mynumber
4	No Zipcode	city Leominster phones

		areacode	339
		number	4120211
		type	work
		areacode	339
		number	8694021
		type	work
		areacode	339
		number	1205678
		type	home
			null
		areacode	305
		number	8064321
		type	home
		state	MA
		street	364 Mulberry Street
6	No Address	myarray	
		mynumber	

7 rows returned

Finally, it is possible to nest search case expressions. Our sample data also has a spurious null in the phones array (see id 4). We can report that in the following way (output is modified slightly to fit in the space allowed):

```
sql-> SELECT id,
CASE
  WHEN EXISTS j.person.address
  THEN
    CASE
      WHEN EXISTS j.person.address.zipcode
      THEN
        CASE
          WHEN j.person.address.phones[] =any null
          THEN "Zipcode exists but null in the phones array"
          ELSE j.person.address.zipcode
        END
      WHEN j.person.address.phones[] =any null
      THEN "No zipcode and null in phones array"
      ELSE "No zipcode"
    END
  ELSE "No Address"
END,
CASE
  WHEN NOT EXISTS j.person.address
  THEN j.person.keys()
  ELSE j.person.address
END
FROM JSONPersons j;
```

id	Column_2	Column_3	
3	No zipcode	city Middleburg	
		phones	
		areacode	305
		number	1234079
		type	work

		areacode number type state street	305 2066401 home FL 187 Aspen Drive
2	53511	city phones areacode number type state street zipcode	Beloit 339 1684972 home WI 187 Hill Street 53511
5	No zipcode	city phones areacode number type areacode number type areacode number type state street	Monroe Township 201 3213267 work 201 8765421 work 339 3414578 home NJ 427 Linden Avenue
1	37013	city phones areacode number type state street zipcode	Antioch 423 8634379 home TN 150 Route 2 37013
7	No Address	myarray mynumber	
4	No zipcode and null in phones array	city phones areacode number type areacode number type areacode number type areacode number	Leominster 339 4120211 work 339 8694021 work 339 1205678 home null 305 8064321

		type	home
		state	MA
		street	364 Mulberry Street
6	No Address	myarray	
		mynumber	

7 rows returned

5

Working With Indexes

The SQL for Oracle NoSQL Database query processor can detect which of the existing indexes on a table can be used to optimize the execution of a query. This chapter provides a brief examples-based introduction to index creation, and queries using indexes. For a more detailed description of index creation and usage, see the *SQL for Oracle NoSQL Database Specification*.

To make it possible to fit the example output on the page, the examples in this chapter use mode `LINE`.

Basic Indexing

This section builds on the examples that you began in [Working with complex data](#).

```
sql-> mode LINE
Query output mode is LINE
sql-> create index idx_income on Persons (income);
Statement completed successfully
sql-> create index idx_age on Persons (age);
Statement completed successfully
sql-> SELECT * from Persons
WHERE income > 10000000 and age < 40;
```

> Row 0

id	3	
firstname	John	
lastname	Morgan	
age	38	
income	100000000	
lastLogin	2016-11-29T08:21:35.4971	
address	street city state zipcode phones	187 Aspen Drive Middleburg FL NULL
	type areacode number	work 305 1234079
	type areacode number	home 305 2066401
connections	1 4	

	2	
expenses	food	2000
	gas	10
	travel	700

1 row returned

Using Index Hints

In the previous section, both indexes are applicable. For index `idx_income`, the query condition `income > 10000000` can be used as the starting point for an index scan that will retrieve only the index entries and associated table rows that satisfy this condition. Similarly, for index `idx_age`, the condition `age < 40` can be used as the stopping point for the index scan. SQL for Oracle NoSQL Database has no way of knowing which of the 2 predicates is more selective, and it assigns the same "value" to each index, eventually picking the one whose name is first alphabetically. In the previous example, `idx_age` was used. To choose the `idx_income` index instead, the query should be written with an index hint:

```
sql-> SELECT /*+ FORCE_INDEX(Persons idx_income) */ * from Persons
WHERE income > 10000000 and age < 40;
```

> Row 0

id	3	
firstname	John	
lastname	Morgan	
age	38	
income	100000000	
lastLogin	2016-11-29T08:21:35.4971	
address	street	187 Aspen Drive
	city	Middleburg
	state	FL
	zipcode	NULL
	phones	
	type	work
	areacode	305
	number	1234079
	type	home
	areacode	305
	number	2066401
connections	1	
	4	
	2	
expenses	food	2000
	gas	10
	travel	700

1 row returned

As shown above, hints are written as a special kind of comment that must be placed immediately after the SELECT keyword. What distinguishes a hint from a regular comment is the "+" character immediately after (without any space) the opening "/*".

Complex Indexes

The following example demonstrates indexing of multiple table fields, indexing of nested fields, and the use of "filtering" predicates during index scans.

```
sql-> create index idx_state_city_income on
Persons (address.state, address.city, income);
Statement completed successfully
sql-> SELECT * from Persons p WHERE p.address.state = "MA"
and income > 79000;
```

> Row 0

id	4	
firstname	Peter	
lastname	Smith	
age	38	
income	80000	
lastLogin	2016-10-19T09:18:05.5555	
address	street	364 Mulberry Street
	city	Leominster
	state	MA
	zipcode	NULL
	phones	
	type	work
	areacode	339
	number	4120211
	type	work
	areacode	339
	number	8694021
	type	home
	areacode	339
	number	1205678
	type	home
	areacode	305
	number	8064321
connections	3	
	5	
	1	
	2	
expenses	books	240
	clothes	2000

	food	6000
	shoes	1200

1 row returned

Index `idx_state_city_income` is applicable to the above query. Specifically, the `state = "MA"` condition can be used to establish the boundaries of the index scan (only index entries whose first field is "MA" will be scanned). Further, during the index scan, the income condition can be used as a "filtering" condition, to skip index entries whose third field is less or equal to 79000. As a result, only rows that satisfy both conditions are retrieved from the table.

Multi-Key Indexes

A multi-key index indexes all the elements of an array, or all the elements and/or all the keys of a map. For such indexes, for each table row, the index contains as many entries as the number of elements/entries in the array/map that is being indexed. Only one array/map may be indexed.

```
sql-> create index idx_areacode on
Persons (address.phones[].areacode);
Statement completed successfully
sql-> SELECT * FROM Persons p WHERE
p.address.phones.areacode =any 339;
```

> Row 0

id	2
firstname	John
lastname	Anderson
age	35
income	100000
lastLogin	2016-11-28T13:01:11.2088
address	street 187 Hill Street
	city Beloit
	state WI
	zipcode 53511
	phones
	type home
	areacode 339
	number 1684972
connections	1
	3
expenses	books 100
	food 1700
	travel 2100

> Row 1

id	4	
firstname	Peter	
lastname	Smith	
age	38	
income	80000	
lastLogin	2016-10-19T09:18:05.5555	
address	street	364 Mulberry Street
	city	Leominster
	state	MA
	zipcode	NULL
	phones	
	type	work
	areacode	339
	number	4120211
	type	work
	areacode	339
	number	8694021
	type	home
	areacode	339
	number	1205678
	type	home
	areacode	305
	number	8064321
connections	3	
	5	
	1	
	2	
expenses	books	240
	clothes	2000
	food	6000
	shoes	1200

> Row 2

id	5	
firstname	Dana	
lastname	Scully	
age	47	
income	400000	
lastLogin	2016-11-08T09:16:46.3929	
address	street	427 Linden Avenue
	city	Monroe Township
	state	NJ

	zipcode	NULL
	phones	
	type	work
	areacode	201
	number	3213267
	type	work
	areacode	201
	number	8765421
	type	home
	areacode	339
	number	3414578
connections	2	
	4	
	1	
	3	
expenses	clothes	1500
	food	900
	shoes	1000

3 rows returned

In the above example, a multi-key index is created on all the area codes in the Persons table, mapping each area code to the persons that have a phone number with that area code. The query is looking for persons who have a phone number with area code 339. The index is applicable to the query and so the key 339 will be searched for in the index and all the associated table rows will be retrieved.

```
sql-> create index idx_expenses on
Persons (expenses.keys(), expenses.values());
Statement completed successfully
sql-> SELECT * FROM Persons p WHERE p.expenses.food > 1000;
```

> Row 0

id	2	
firstname	John	
lastname	Anderson	
age	35	
income	100000	
lastLogin	2016-11-28T13:01:11.2088	
address	street	187 Hill Street
	city	Beloit
	state	WI
	zipcode	53511
	phones	
	type	home
	areacode	339
	number	1684972

connections	1	
	3	
expenses	books	100
	food	1700
	travel	2100

> Row 1

id	3	
firstname	John	
lastname	Morgan	
age	38	
income	100000000	
lastLogin	2016-11-29T08:21:35.4971	
address	street	187 Aspen Drive
	city	Middleburg
	state	FL
	zipcode	NULL
	phones	
	type	work
	areacode	305
	number	1234079
	type	home
	areacode	305
	number	2066401
connections	1	
	4	
	2	
expenses	food	2000
	gas	10
	travel	700

> Row 2

id	4	
firstname	Peter	
lastname	Smith	
age	38	
income	80000	
lastLogin	2016-10-19T09:18:05.5555	
address	street	364 Mulberry Street
	city	Leominster
	state	MA

	zipcode	NULL
	phones	
	type	work
	areacode	339
	number	4120211
	type	work
	areacode	339
	number	8694021
	type	home
	areacode	339
	number	1205678
	type	home
	areacode	305
	number	8064321

connections	3	
	5	
	1	
	2	

expenses	books	240
	clothes	2000
	food	6000
	shoes	1200

3 rows returned

In the above example, a multi-key index is created on all the expenses entries in the Persons table, mapping each category C and each amount A associated with that category to the persons that have an entry (C, A) in their expenses map. The query is looking for persons who spent more than 1000 on food. The index is applicable to the query and so only the index entries whose first field (the map key) is equal to "food" and second key (the amount) is greater than 1000 will be scanned and the associated rows retrieved.

Indexing JSON Data

An index is a JSON index if it indexes at least one field that is contained inside JSON data.

Because JSON is schema-less, it is possible for JSON data to differ in type across table rows. However, when indexing JSON data, the data type must be consistent across table rows or the index creation will fail. Further, once one or more JSON indexes have been created, any attempt to write data of an incorrect type will fail.

With the exception of the previous restriction, indexing JSON data and working with JSON indexes behaves in much the same way as indexing non-JSON data. To create the index, specify a path to the JSON field using dot notation. You must also specify the data's type, using the `AS` keyword.

The following examples are built on the examples shown in [Working with JSON](#).

```
sql-> create index idx_json_income on JSONPersons (person.income
as integer);
Statement completed successfully
```

```
sql-> create index idx_json_age on JSONPersons (person.age as integer);
Statement completed successfully
sql->
```

You can then run a query in the normal way, and the index `idx_json_income` will be automatically used. But as shown at the beginning of this chapter ([Basic Indexing](#)), the query processor will not know which index to use. To require the use of a particular index provide an index hint as normal:

```
sql-> SELECT /*+ FORCE_INDEX(JSONPersons idx_json_income) */ *
from JSONPersons j WHERE j.person.income > 10000000 and
j.person.age < 40;
```

```
> Row 0
```

```
+-----+-----+
| id      | 3      |
+-----+-----+
| person  | address | |
|         | city    | Middleburg |
|         | phones  |            |
|         |   areacode | 305      |
|         |   number  | 1234079  |
|         |   type    | work     |
|         |         |          |
|         |   areacode | 305      |
|         |   number  | 2066401  |
|         |   type    | home     |
|         |   state   | FL       |
|         |   street  | 187 Aspen Drive |
|         | age      | 38       |
|         | connections |          |
|         |         | 1        |
|         |         | 4        |
|         |         | 2        |
|         | expenses |          |
|         |   food   | 2000     |
|         |   gas    | 10       |
|         |   travel | 700      |
|         |   firstname | John    |
|         |   income | 100000000 |
|         |   lastLogin | 2016-11-29T08:21:35.4971 |
|         |   lastname | Morgan  |
+-----+-----+
```

```
1 row returned
```

```
sql->
```

Finally, when creating a multi-key index on a JSON map, a type must not be given for the `.keys()` expression. This is because the type will always be `String`. However, a type declaration is required for the `.values()` expression:

```
sql-> create index idx_json_expenses on JSONPersons
(person.expenses.keys(), person.expenses.values() as integer);
Statement completed successfully
sql-> SELECT * FROM JSONPersons j WHERE j.person.expenses.food > 1000;
```

```
> Row 0
```

```
+-----+-----+
| id      | 2      |
+-----+-----+
| person  | address |
+-----+-----+
```

	city	Beloit
	phones	
	areacode	339
	number	1684972
	type	home
	state	WI
	street	187 Hill Street
	zipcode	53511
	age	35
	connections	
		1
		3
	expenses	
	books	100
	food	1700
	travel	2100
	firstname	John
	income	100000
	lastLogin	2016-11-28T13:01:11.2088
	lastname	Anderson

> Row 1

id	3	
person	address	
	city	Middleburg
	phones	
	areacode	305
	number	1234079
	type	work
	areacode	305
	number	2066401
	type	home
	state	FL
	street	187 Aspen Drive
	age	38
	connections	
		1
		4
		2
	expenses	
	food	2000
	gas	10
	travel	700
	firstname	John
	income	100000000
	lastLogin	2016-11-29T08:21:35.4971
	lastname	Morgan

> Row 2

id	4	
person	address	
	city	Leominster
	phones	
	areacode	339

6

Modifying Table Rows using UPDATE Statements

This chapter provides examples of how to update table rows using SQL for Oracle NoSQL Database UPDATE statements. These are an efficient way to update table row data, because UPDATE statements make *server-side updates* directly, without requiring a Read/Modify/Write update cycle.

 **Note:**

You can use UPDATE statements to update only an existing row. You cannot use UPDATE to either create new rows, or delete existing rows. An UPDATE statement can modify only a single row at a time.

The UPDATE statement syntax is described in [SQL UPDATE Statements](#)

Example Data

This chapter's examples uses the data loaded by the `SQLJSONExamples` script, which can be found in the `Examples` download package. For details on using this script, the sample data it loads, and the `Examples` download, see [SQLJSONExamples Script](#).

Changing Field Values

In the simplest case, you can change the value of a field using the Update Statement SET clause. The JSON example data set has a row which contains just an array and an integer. This is row ID 6:

```
sql-> mode column
Query output mode is COLUMN
sql-> SELECT * from JSONPersons j WHERE j.id = 6;
```

id	person
6	myarray 1 2 3 4 mynumber 5

1 row returned

You can change the value of `mynumber` in that row using the following statement:

```
sql-> UPDATE JSONPersons j
      SET j.person.mynumber = 100
      WHERE j.id = 6;
```

```
+-----+
| Column_1 |
+-----+
|         1 |
+-----+
```

1 row returned

```
sql-> SELECT * from JSONPersons j WHERE j.id = 6;
```

```
+-----+-----+
| id |      person      |
+-----+-----+
|  6 | myarray          | |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   | mynumber | 100 |
+-----+-----+
```

1 row returned

In the previous example, the results returned by the Update statement was not very informative, so we were required to reissue the Select statement in order to view the results of the update. You can avoid that by using a RETURNING clause. This functions exactly like a Select statement:

```
sql-> UPDATE JSONPersons j
      SET j.person.mynumber = 200
      WHERE j.id = 6
      RETURNING *;
```

```
+-----+-----+
| id |      person      |
+-----+-----+
|  6 | myarray          | |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   |                  |
|   | mynumber | 200 |
+-----+-----+
```

1 row returned

```
sql->
```

You can further limit and customize the displayed results in the same way that you can do so using a SELECT statement:

```
sql-> UPDATE JSONPersons j
      SET j.person.mynumber = 300
      WHERE j.id = 6
      RETURNING id, j.person.mynumber AS MyNumber;
```

```
+-----+-----+
| id | MyNumber |
+-----+-----+
|  6 | 300      |
+-----+-----+
```

```
1 row returned
sql->
```

It is normally possible to update the value of a non-JSON field using the SET clause. However, you cannot change a field if it is a primary key. For example:

```
sql-> UPDATE JSONPersons j
      SET j.id = 1000
      WHERE j.id = 6
      RETURNING *;
Error handling command UPDATE JSONPersons j
SET j.id = 1000
WHERE j.id = 6
RETURNING *: Error: at (2, 4) Cannot update a primary key column
Usage:

Unknown statement

sql->
```

Modifying Array Values

You use the Update statement ADD clause to add elements into an array. You use a SET clause to change the value of an existing array element. And you use a REMOVE clause to remove elements from an array.

Adding Elements to an Array

The ADD clause requires you to identify the array position that you want to operate on, followed by the value you want to set to that position in the array. If the index value that you set is 0 or a negative number, the value that you specify is inserted at the beginning of the array.

If you do not provide an index position, the array value that you specify is appended to the end of the array.

```
sql-> SELECT * from JSONPersons j WHERE j.id = 6;
```

id	person
6	myarray
	1
	2
	3
	4
	mynumber 300

```
1 row returned
```

```
sql-> UPDATE JSONPersons j
      ADD j.person.myarray 0 50,
      ADD j.person.myarray 100
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	myarray
	50


```
1 row returned
sql->
```

If you provide an array position that is between 0 and the array's size, then the value you specify will be inserted into the array *before* the specified position. To determine the correct position, start counting from 0:

```
UPDATE JSONPersons j
  ADD j.person.myarray 3 250
  WHERE j.id = 6
  RETURNING *;
```

id	person
6	myarray
	50
	1
	2
	250
	3
	4
	100
	400
	66
	77
	88
	mynumber 300

```
1 row returned
sql->
```

Changing an Existing Element in an Array

To change an existing value in an array, use the SET clause and identify the value's position using []. To determine the value's position, start counting from 0:

```
sql-> UPDATE JSONPersons j
      SET j.person.myarray[3] = 1000
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	myarray
	50
	1
	2
	1000
	3
	4
	100
	400
	66
	77
	88
	mynumber 300

```
1 row returned
sql->
```

Removing Elements from Arrays

To remove an existing element from an array, use the REMOVE clause. To do this, you must identify the position of the element in the array that you want to remove. To determine the value's position, start counting from 0:

```
sql-> UPDATE JSONPersons j
      REMOVE j.person.myarray[3]
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	myarray
	50
	1
	2
	3
	4
	100
	400
	66
	77
	88
	mynumber 300

```
1 row returned
sql->
```

It is possible for the array position to be identified by an expression. For example, in our sample data, some records include an array of phone numbers, and some of those phone numbers include a work number:

```
sql-> SELECT * FROM JSONPersons j WHERE j.id = 3;
```

id	person
3	address
	city Middleburg
	phones
	areacode 305
	number 1234079
	type work
	areacode 305
	number 2066401
	type home
	state FL
	street 187 Aspen Drive
	age 38
	connections
	1
	4
	2
	expenses
	food 2000

gas	10
travel	700
firstname	John
income	100000000
lastLogin	2016-11-29T08:21:35.4971
lastname	Morgan

1 row returned
sql->

We can remove the work number from the array in one of two ways. First, we can directly specify its position in the array (position 0), but that only removes a single element at a time. If we want to remove all the work numbers, we can do it by using the \$element variable. To illustrate, we first add another work number to the array:

```
sql-> UPDATE JSONPersons j
      ADD j.person.address.phones 0
      {"type":"work", "areacode":415, "number":9998877}
      WHERE j.id = 3
      RETURNING *;
```

id	person	
3	address	
	city	Middleburg
	phones	
	areacode	415
	number	9998877
	type	work
	areacode	305
	number	1234079
	type	work
	areacode	305
	number	2066401
	type	home
	state	FL
	street	187 Aspen Drive
	age	38
	connections	
		1
		4
		2
	expenses	
	food	2000
	gas	10
	travel	700
	firstname	John
	income	100000000
	lastLogin	2016-11-29T08:21:35.4971
	lastname	Morgan

1 row returned
sql->

Now we can remove all the work numbers as follows:

```
sql-> UPDATE JSONPersons j
      REMOVE j.person.address.phones[$element.type = "work"]
      WHERE j.id = 3
      RETURNING *;
```

id	person
3	address city Middleburg phones areacode 305 number 2066401 type home state FL street 187 Aspen Drive age 38 connections 1 4 2 expenses food 2000 gas 10 travel 700 firstname John income 100000000 lastLogin 2016-11-29T08:21:35.4971 lastname Morgan

```
1 row returned
sql->
```

Modifying Map Values

To write a new field to a map, use the PUT clause. You can also use the PUT clause to change an existing map value. To remove a map field, use the REMOVE clause.

For example, consider the following two rows from our sample data:

```
sql-> SELECT * FROM JSONPersons j WHERE j.id = 6 OR j.id = 3;
```

id	person
3	address city Middleburg phones areacode 305 number 2066401 type home state FL street 187 Aspen Drive age 38 connections 1 4 2 expenses food 2000 gas 10

	travel	700
	firstname	John
	income	100000000
	lastLogin	2016-11-29T08:21:35.4971
	lastname	Morgan

6	myarray	
		50
		1
		2
		3
		4
		100
		400
		66
		77
		88
	mynumber	300

2 rows returned
sql->

These two rows look nothing alike. Row 3 contains information about a person, while row 6 contains, essentially, random data. This is possible because the `person` column is of type JSON, which is not strongly typed. But because we interact with JSON columns as if they are maps, we can fix row 6 by modifying it as a map.

Removing Elements from a Map

To begin, we remove the two existing elements from row six (`myarray` and `mynumber`). We do this with a single UPDATE statement, which allows us to execute multiple update clauses so long as they are comma-separated:

```
sql-> UPDATE JSONPersons j
      REMOVE j.person.myarray,
      REMOVE j.person.mynumber
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	

1 row returned
sql->

Adding Elements to a Map

Next, we add person data to this table row. We could do this with a single UPDATE statement by specifying the entire map with a single PUT clause, but for illustration purposes we do this in multiple steps.

To begin, we specify the person's name. Here, we use a single PUT clause that specifies a map with multiple elements:

```
sql-> UPDATE JSONPersons j
      PUT j.person {"firstname" : "Wendy",
```

```

        "lastname" : "Purvis"}
WHERE j.id = 6
RETURNING *;
+-----+-----+
| id |          person          |
+-----+-----+
| 6 | firstname | Wendy |
|   | lastname  | Purvis |
+-----+-----+

```

1 row returned
sql->

Next, we specify the age, connections, expenses, income, and lastLogin fields using multiple PUT clauses on a single UPDATE statement:

```

sql-> UPDATE JSONPersons j
      PUT j.person {"age" : 43},
      PUT j.person {"connections" : [2,3]},
      PUT j.person {"expenses" : {"food" : 1100,
                                  "books" : 210,
                                  "travel" : 50}},
      PUT j.person {"income" : 80000},
      PUT j.person {"lastLogin" : "2017-06-29T16:12:35.0285"}
WHERE j.id = 6
RETURNING *;

```

```

+-----+-----+
| id |          person          |
+-----+-----+
| 6 | age          | 43 |
|   | connections  |    |
|   |              | 2  |
|   |              | 3  |
|   | expenses    |    |
|   |   books    | 210 |
|   |   food     | 1100 |
|   |   travel   | 50  |
|   | firstname  | Wendy |
|   | income     | 80000 |
|   | lastLogin  | 2017-06-29T16:12:35.0285 |
|   | lastname   | Purvis |
+-----+-----+

```

1 row returned
sql->

We still need an address. Again, we could do this with a single PUT clause, but for illustration purposes we will use multiple clauses. Our first PUT creates the `address` element, which uses a map as a value. Our second PUT adds elements to the `address` map:

```

sql-> UPDATE JSONPersons j
      PUT j.person {"address" : {"street" : "479 South Way Dr"}},
      PUT j.person.address {"city" : "St. Petersburg",
                            "state" : "FL"}
WHERE j.id = 6
RETURNING *;

```

```

+-----+-----+
| id |          person          |
+-----+-----+
| 6 | address                |
+-----+-----+

```

city	St. Petersburg
state	FL
street	479 South Way Dr
age	43
connections	2
	3
expenses	
books	210
food	1100
travel	50
firstname	Wendy
income	80000
lastLogin	2017-06-29T16:12:35.0285
lastname	Purvis

1 row returned
sql->

Finally, we provide phone numbers for this person. These are specified as an array of maps:

```
sql-> UPDATE JSONPersons j
      PUT j.person.address {"phones" :
        [{"type":"work", "areacode":727, "number":8284321},
         {"type":"home", "areacode":727, "number":5710076},
         {"type":"mobile", "areacode":727, "number":8913080}
        ]
      }
      WHERE j.id = 6
      RETURNING *;
```

id	person																																												
6	address <table border="1"> <tr><td>city</td><td>St. Petersburg</td></tr> <tr><td>phones</td><td></td></tr> <tr><td> areacode</td><td>727</td></tr> <tr><td> number</td><td>8284321</td></tr> <tr><td> type</td><td>work</td></tr> <tr><td> areacode</td><td>727</td></tr> <tr><td> number</td><td>5710076</td></tr> <tr><td> type</td><td>home</td></tr> <tr><td> areacode</td><td>727</td></tr> <tr><td> number</td><td>8913080</td></tr> <tr><td> type</td><td>mobile</td></tr> <tr><td>state</td><td>FL</td></tr> <tr><td>street</td><td>479 South Way Dr</td></tr> <tr><td>age</td><td>43</td></tr> <tr><td>connections</td><td>2</td></tr> <tr><td></td><td>3</td></tr> <tr><td>expenses</td><td></td></tr> <tr><td> books</td><td>210</td></tr> <tr><td> food</td><td>1100</td></tr> <tr><td> travel</td><td>50</td></tr> <tr><td>firstname</td><td>Wendy</td></tr> <tr><td>income</td><td>80000</td></tr> </table>	city	St. Petersburg	phones		areacode	727	number	8284321	type	work	areacode	727	number	5710076	type	home	areacode	727	number	8913080	type	mobile	state	FL	street	479 South Way Dr	age	43	connections	2		3	expenses		books	210	food	1100	travel	50	firstname	Wendy	income	80000
city	St. Petersburg																																												
phones																																													
areacode	727																																												
number	8284321																																												
type	work																																												
areacode	727																																												
number	5710076																																												
type	home																																												
areacode	727																																												
number	8913080																																												
type	mobile																																												
state	FL																																												
street	479 South Way Dr																																												
age	43																																												
connections	2																																												
	3																																												
expenses																																													
books	210																																												
food	1100																																												
travel	50																																												
firstname	Wendy																																												
income	80000																																												


```

|      | lastLogin      | 2017-06-29T16:12:35.0285 |
|      | lastname      | Purvis                    |
+-----+

```

```

1 row returned
sql->

```

Updating Existing Map Elements

To update an existing element in a map, you can use the PUT clause in exactly the same way as you add a new element to map. For example, to update the lastLogin time:

```

sql-> UPDATE JSONPersons j
      PUT j.person {"lastLogin" : "2017-06-29T20:36:04.9661"}
      WHERE j.id = 6
      RETURNING *;

```

```

+-----+
| id |                person                |
+-----+
| 6  | address                               | |
|    |   city                               | St. Petersburg |
|    |   phones                             |               |
|    |     areacode                         | 727           |
|    |     number                           | 8284321       |
|    |     type                             | work          |
|    |                                       |               |
|    |     areacode                         | 727           |
|    |     number                           | 5710076       |
|    |     type                             | home          |
|    |                                       |               |
|    |     areacode                         | 727           |
|    |     number                           | 8913080       |
|    |     type                             | mobile        |
|    |   state                             | FL            |
|    |   street                             | 479 South Way Dr |
|    | age                                  | 43            |
|    | connections                          |               |
|    |                                       | 2             |
|    |                                       | 3             |
|    | expenses                             |               |
|    |   books                              | 210           |
|    |   food                               | 1100          |
|    |   travel                             | 50            |
|    |   firstname                          | Wendy         |
|    |   income                             | 80000         |
|    |   lastLogin                          | 2017-06-29T20:36:04.9661 |
|    |   lastname                           | Purvis        |
+-----+

```

```

1 row returned
sql->

```

Alternatively, use a SET clause:

```

sql-> UPDATE JSONPersons j
      SET j.person.lastLogin = "2017-06-29T20:38:56.2751"
      WHERE j.id = 6
      RETURNING *;

```

```

+-----+

```

id	person
6	address
	city St. Petersburg
	phones
	areacode 727
	number 8284321
	type work
	areacode 727
	number 5710076
	type home
	areacode 727
	number 8913080
	type mobile
	state FL
	street 479 South Way Dr
	age 43
	connections
	2
	3
	expenses
	books 210
	food 1100
	travel 50
	firstname Wendy
	income 80000
	lastLogin 2017-06-29T20:38:56.2751
	lastname Purvis

1 row returned
sql->

If you want to set the timestamp to the current time, use the `current_time()` built-in function (see [Time Functions](#)):

```
sql-> UPDATE JSONPersons j
      SET j.person.lastLogin = cast(current_time() AS String)
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	address
	city St. Petersburg
	phones
	areacode 727
	number 8284321
	type work
	areacode 727
	number 5710076
	type home
	areacode 727
	number 8913080
	type mobile
	state FL
	street 479 South Way Dr

age	43
connections	2
	3
expenses	
books	210
food	1100
travel	50
firstname	Wendy
income	80000
lastLogin	2017-06-29T04:40:15.917
lastname	Purvis

1 row returned
sql->

If an element in the map is an array, you can modify it in the same way as you would any array. For example:

```
sql-> UPDATE JSONPersons j
      ADD j.person.connections seq_concat(1, 4)
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	address
	city St. Petersburg
	phones
	areacode 727
	number 8284321
	type work
	areacode 727
	number 5710076
	type home
	areacode 727
	number 8913080
	type mobile
	state FL
	street 479 South Way Dr
	age 43
	connections
	2
	3
	1
	4
	expenses
	books 210
	food 1100
	travel 50
	firstname Wendy
	income 80000
	lastLogin 2017-06-29T04:40:15.917
	lastname Purvis

1 row returned

If you are unsure of an element being an array or a map, you can use both ADD and PUT within the same UPDATE statement. For example:

```
sql-> UPDATE JSONPersons j
      ADD j.person.connections seq_concat(5, 7),
      PUT j.person.connections seq_concat(5, 7)
      WHERE j.id = 6
      RETURNING *;
```

id	person
6	address city St. Petersburg phones areacode 727 number 8284321 type work areacode 727 number 5710076 type home areacode 727 number 8913080 type mobile state FL street 479 South Way Dr age 43 connections 2 3 1 4 5 7 expenses books 210 food 1100 travel 50 firstname Wendy income 80000 lastLogin 2017-06-29T04:40:15.917 lastname Purvis

1 row returned

If the element is an array, the ADD gets applied and the PUT is a noop. If it is a map, then the PUT gets applied and ADD is a noop. In this example, since the element is an array, the ADD gets applied.

Managing Time to Live Values

Time to Live (TTL) values indicate how long data can exist in a table before it expires. Expired data can no longer be returned as part of a query.

Default TTL values can be set on either a table-level or a row level when the table is first defined. Using UPDATE statements, you can change the TTL value for a single row.

You can see a row's TTL value using the `remaining_hours()`, `remaining_days()` or `expiration_time()` built-in functions. These TTL functions require a row as input. We accomplish this by using the `$` as part of the table alias. This causes the table alias to function as a row variable.

```
sql-> SELECT remaining_days($j) AS Expires
      FROM JSONPersons $j WHERE id = 6;
+-----+
| Expires |
+-----+
|      -1 |
+-----+
```

1 row returned
sql->

The previous query returns `-1`. This means that the row has no expiration time. We can specify an expiration time for the row by using an `UPDATE` statement with a `set TTL` clause. This clause computes a new TTL by specifying an offset from the current expiration time. If the row never expires, then the current expiration time is `1970-01-01T00:00:00.000`. The value you provide to `set TTL` must specify units of either `HOURS` or `DAYS`.

```
sql-> UPDATE JSONPersons $j
      SET TTL 1 DAYS
      WHERE id = 6
      RETURNING remaining_days($j) AS Expires;
+-----+
| Expires |
+-----+
|        1 |
+-----+
```

1 row returned
sql->

To see the new expiration time, we can use the built-in `expiration_time()` function. Because we specified an expiration time based on a day boundary, the row expires at midnight of the following day (expiration rounds up):

```
sql-> SELECT current_time() AS Now,
      expiration_time($j) AS Expires
      FROM JSONPersons $j WHERE id = 6;
+-----+-----+
|          Now          | Expires          |
+-----+-----+
| 2017-07-03T21:56:47.778 | 2017-07-05T00:00:00.000 |
+-----+-----+
```

1 row returned
sql->

To turn off the TTL so that the row will never expire, specify a negative value, using either `HOURS` or `DAYS` as the unit:

```
sql-> UPDATE JSONPersons $j
      SET TTL -1 DAYS
      WHERE id = 6
      RETURNING remaining_days($j) AS Expires;
+-----+
```

```

| Expires |
+-----+
|      0 |
+-----+

```

1 row returned
sql->

Notice that the RETURNING clause provides a value of 0 days. This indicates that the row will never expire. Further, if we look at the remaining_days() using a SELECT statement, we will once again see a negative value, indicating that the row never expires:

```

sql-> SELECT remaining_days($j) AS Expires
      FROM JSONPersons $j WHERE id = 6;
+-----+
| Expires |
+-----+
|     -1 |
+-----+

```

1 row returned
sql->

Avoiding the Read-Modify-Write Cycle

An important aspect of UPDATE Statements is that you do not have to read a value in order to update it. Instead, you can blindly modify a value directly in the store without ever retrieving (reading) it. To do this, you refer to the value you want to modify using the \$ variable.

For example, we have a row in JSONPersons that looks like this:

```

sql-> SELECT * FROM JSONPersons WHERE id=6;
+-----+-----+-----+
| id |          person          |
+-----+-----+-----+
| 6 | address                  |
|   |   city                   | St. Petersburg
|   |   phones                 |
|   |     areacode             | 727
|   |     number               | 8284321
|   |     type                 | work
|   |
|   |     areacode             | 727
|   |     number               | 5710076
|   |     type                 | home
|   |
|   |     areacode             | 727
|   |     number               | 8913080
|   |     type                 | mobile
|   |   state                 | FL
|   |   street                 | 479 South Way Dr
|   |   age                   | 43
|   | connections              |
|   |                           | 2
|   |                           | 3
|   |                           | 1
|   |                           | 4
|   | expenses                 |
+-----+-----+-----+

```

books	210
food	1100
travel	50
firstname	Wendy
income	80000
lastLogin	2017-07-25T22:50:06.482
lastname	Purvis

1 row returned

We can blindly update the value of the `person.expenses.books` field by referencing `$`. In the following statement, no read is performed on the store. Instead, the write operation is performed directly at the store.

```
sql-> UPDATE JSONPersons j
->     SET j.person.expenses.books = $ + 100
->     WHERE id = 6;
```

NumRowsUpdated	1
----------------	---

1 row returned

To see that the books expenses value has indeed been incremented by 100, we perform a second `SELECT` statement.

```
sql-> SELECT * FROM JSONPersons WHERE id=6;
```

id	person
6	address
	city St. Petersburg
	phones
	areacode 727
	number 8284321
	type work
	areacode 727
	number 5710076
	type home
	areacode 727
	number 8913080
	type mobile
	state FL
	street 479 South Way Dr
	age 43
	connections
	2
	3
	1
	4
	expenses
	books 310
	food 1100
	travel 50
	firstname Wendy
	income 80000

Part II

Language Definition

This part provides a textual overview of SQL for Oracle NoSQL Database. While some examples are provided to illustrate concepts, more thorough examples are provided in [Introductory Examples](#).

7

The SQL for Oracle NoSQL Database Data Model

This chapter gives an overview of the data model for SQL for Oracle NoSQL Database. For a more detailed description of the data model see the *SQL for Oracle NoSQL Database Specification*.

Example Data

The language definition portion of this document frequently provides examples to illustrate the concepts. The following table definition is used by those examples:

```
CREATE TABLE Users (  
    id INTEGER,  
    firstName STRING,  
    lastName STRING,  
    otherNames ARRAY(RECORD(first STRING, last STRING)),  
    age INTEGER,  
    income INTEGER,  
    address JSON,  
    connections ARRAY(INTEGER),  
    expenses MAP(INTEGER),  
    moveDate timestamp(4),  
    PRIMARY KEY (id)  
)
```

The rows of the Users table defined above represent information about users. For each such user, the “connections” field is an array containing ids of other users that this user is connected with. The ids in the array are sorted by some measure of the strength of the connection.

The “expenses” column is a maps expense categories (like “housing”, clothes”, “books”, etc) to the amount spent in the associated category. The set of categories may not be known in advance, or it may differ significantly from user to user, or may need to be frequently updated by adding or removing categories for each user. As a result, using a map type for “expenses”, instead of a record type, is the right choice.

Finally, the “address” column has type JSON. A typical value for “address” will be a map representing a json document.

Typical row data for this table will look like this:

```
{  
    "id":1,  
    "firstname":"David",  
    "lastname":"Morrison",  
    "otherNames" : [{"first" : "Dave",  
                    "last" : "Morrison"}],  
    "age":25,  
    "income":100000,  
    "address":{"street":"150 Route 2",  
              "city":"Antioch",
```

```
        "state": "TN",
        "zipcode" : 37013,
        "phones": [{"type": "home", "areacode": 423,
                    "number": 8634379}]
    },
    "connections": [2, 3],
    "expenses": {"food": 1000, "gas": 180},
    "moveDate" : "2016-10-29T18:43:59.8319"
}
```

Data Types and Values

In SQL for Oracle NoSQL Database data is modeled as typed items. A typed item (or simply item) is a value and an associated type that contains the value. A type is a definition of a set of values that are said to belong to (or be instances of) that type.

Values can be atomic or complex. An atomic value is a single, indivisible unit of data. A complex value is a value that contains or consists of other values and provides access to its nested values. Similarly, the types supported by SQL for Oracle NoSQL Database can be characterized as atomic types (containing atomic values only) or complex types (containing complex values only).

The data model supports the following kinds of atomic values and associated data types:

- Integer
4-byte long integer number.
- Long
8-byte long integer number.
- Float
4-byte long real number.
- Double
8-byte long real number.
- Number
All numbers that can be represented by the Java `BigDecimal` data type.
- String
Sequence of unicode characters.
- Boolean
Values are either `true` or `false`.
- Binaries
An uninterpreted sequence of zero or more bytes.
- Enums
Values are symbolic identifiers (tokens). Enums are stored as strings, but are not considered to be strings.
- Timestamp

Represents a point in time as a date and, optionally, a time. See [Timestamp](#) for details.

- JSON null value

Special value that indicates the absence of an actual value within a JSON datatype such as an object, map, or array.

- SQL NULL

Special value that is used to indicate the absence of an actual value, or the fact that a value is unknown or inapplicable.

SQL for Oracle NoSQL Database also supports the following complex types:

- Array

An array is an ordered collection of zero or more items. Normally all elements of an array have the same type. Also, normally arrays cannot contain NULL items. However, arrays of type JSON can contain a mix of JSON datatypes, as well as NULL items.

- Map

An unordered collection of zero or more key-item pairs, where all keys are strings and all the items normally have the same type. Also, normally Maps cannot contain NULL items. However, if the map is of type JSON, then it can contain a mix of datatypes for the items, as well as NULL items.

- Record

An ordered collection of one or more key-item pairs, where all keys are strings and the items associated with different keys may have different types. Also, record items may be NULL.

Another difference between records and maps is that the keys in records are fixed and known in advance (they are part of the record type definition), whereas maps can contain arbitrary keys (the map keys are not part of the map type).

Wildcard Types and JSON Data

The Oracle NoSQL data model includes the following wildcard types:

- Any

All possible values.

- AnyAtomic

All possible atomic values.

- AnyJsonAtomic

All atomic values that are valid JSON values. This is the union of all numeric values, all string values, the 2 boolean values, and the JNULL value.

- JSON

All possible JSON values. The domain set is defined recursively as follows:

1. All AnyJsonAtomic values.
2. All arrays whose elements are included in AnyJsonAtomic values.
3. All maps whose field values are those described in (1) and (2) of this list.

- AnyRecord
All possible record values.

With the exception of JNULL items (which pair the JNULL value with the JSON type), no item can have a wildcard type as its type. Wildcard types should be viewed as abstract types. However, items may have an imprecise type. For example, an item may have `Map(JSON)` as its type, indicating that its value is a map that can store field values of different types, as long as all of these values belong to the JSON type. In fact, `Map(JSON)` is the type that represents all JSON objects (JSON documents), and `Array(JSON)` is the type that represents all JSON arrays.

**Note:**

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

JSON Data

To load JSON data into a table, input is accepted as strings or streams containing JSON text. Oracle NoSQL Database parses the input text internally and maps its constituent pieces to the values and types of the data model described here.

Specifically, when an array is encountered in the input text, an array item is created whose type is `Array(JSON)`. This is done unconditionally, no matter what the actual contents of the array might be. For example, even if the array contains integers only, the array item that will be created will have type `Array(JSON)`. The reason that the array is not created with type `Array(Integer)` is that this would mean that we could never update the array by putting something other than integers.

For the same reason, when a JSON object is encountered in the input text, a map item is created whose type is `Map(JSON)`, unconditionally. When numbers are encountered, they are converted to integer, long, or double items, depending on the actual value of the number (float items are not used for JSON). Finally, strings in the input text are mapped to string items, boolean values are mapped to boolean items, and JSON nulls to JSON null items.

Timestamp

Represents a point in time as a date and, optionally, a time value.

Timestamp values have a precision in fractional seconds that range from 0 to 9. For example, a precision of 0 means that no fractional seconds are stored, 3 means that the timestamp stores milliseconds, and 9 means a precision of nanoseconds. 0 is the minimum precision, and 9 is the maximum.

There is no timezone information stored in timestamp; they are all assumed to be in the UTC timezone.

The number of bytes used to store a timestamp depends on its precision (the on-disk storage varies between 5 and 9 bytes).

This datatype is specified as a string in the following format:

```
"<yyyy>-<mm>-<dd>[T<HH>:<mm>:<ss>[.<SS>]]"
```

where:

- <yyyy> is the four-digit year value (for example, "2016").
- <mm> is the two-digit month value (for example, "08").
- <dd> is the two-digit day value (for example, "01").
- <HH> is the two-digit hour value (for example, "18").
- <mm> is the two-digit minute value.
- <ss> is the two-digit seconds value.
- <SS> is the fractional seconds value. If the value specified here exceeds the precision declared when the table column was defined, then the value specified is rounded off.

You can return the current time as a timestamp using the `current_time()` function. See [Time Functions](#) for details.

Timestamp Functions

The following functions can be used with the timestamp datatype:

- `year(<timestamp>)`
Returns the year for the given timestamp. The returned value is in the range -6383 to 9999. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.
- `month(<timestamp>)`
Returns the month for the given timestamp. The returned value is in the range 1 to 12. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.
- `day(<timestamp>)`
Returns the day for the given timestamp. The returned value is in the range 1 to 31. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.
- `hour(<timestamp>)`
Returns the hour for the given timestamp. The returned value is in the range 0 to 23. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.
- `minute(<timestamp>)`
Returns the minute for the given timestamp. The returned value is in the range 0 to 59. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.
- `second(<timestamp>)`
Returns the second for the given timestamp. The returned value is in the range 0 to 59. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.
- `millisecond(<timestamp>)`

Returns the millisecond for the given timestamp. The returned value is in the range 0 to 999. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.

- `microsecond(<timestamp>)`

Returns the microsecond for the given timestamp. The returned value is in the range 0 to 999999. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.

- `nanosecond(<timestamp>)`

Returns the nanosecond for the given timestamp. The returned value is in the range 0 to 999999999. If the <timestamp> argument is NULL or empty, the result is also NULL or empty.

- `week(<timestamp>)`

Returns the week number within the year. Weeks start on a Sunday, and the first week in the year has a minimum of 1 day. The value returned is in the range 1 to 54. If the argument is NULL or empty, the result is also NULL or empty.

- `isoweek(<timestamp>)`

Returns the week number within the year based on ISO-8601. Weeks start on Monday, and the first week has a minimum of 4 days. The value returned is in the range 0 to 53. If the argument is NULL or empty, the result is also NULL or empty.

- `extract(<unit> from <expr>)`

Extracts temporal fields from a timestamp field. <expr> must return a timestamp. <unit> must be one of YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, MILLISECOND, MICROSECOND, NANOSECOND, WEEK, or ISOWEEK.

See [Working with Timestamps](#) for an example of using these functions.

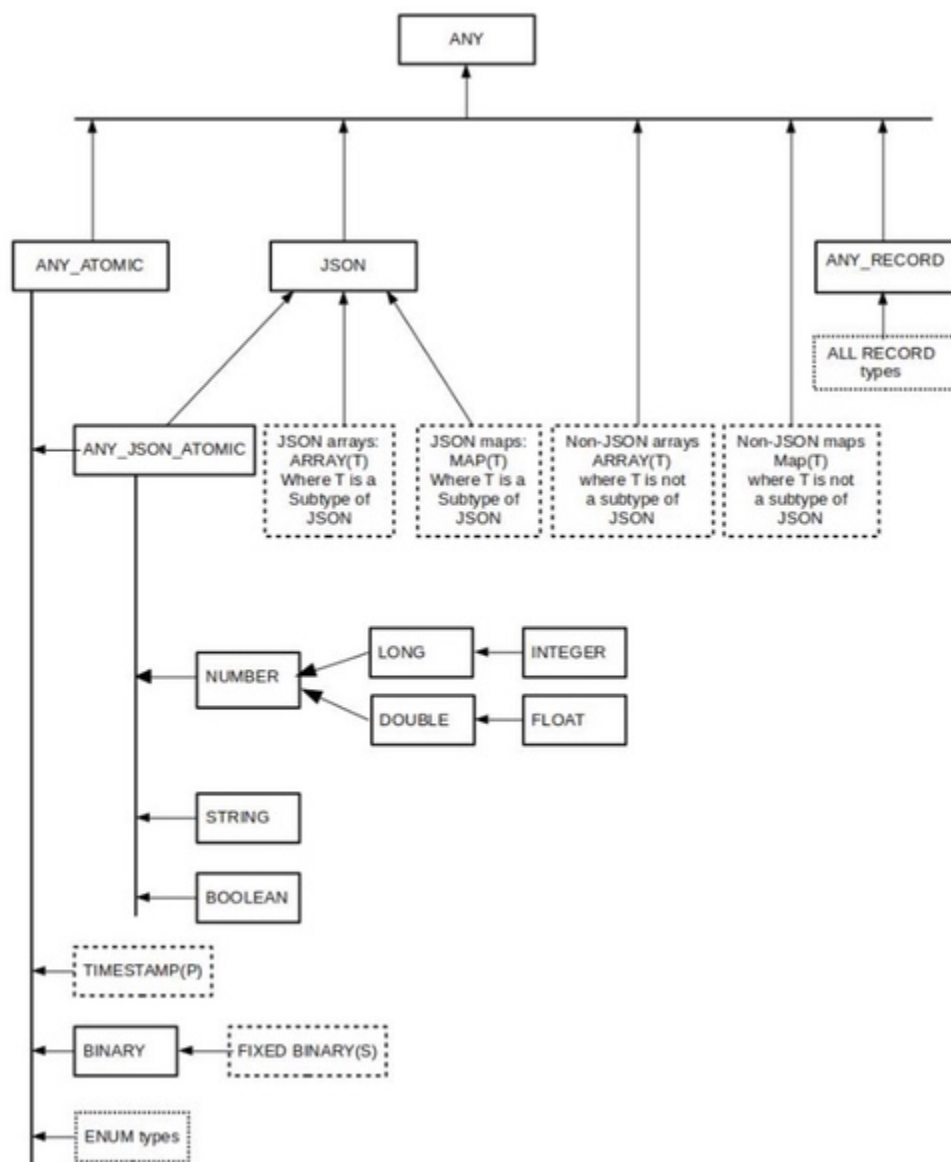
Type Hierarchy

SQL for Oracle NoSQL Database defines a subtype-supertype relationship among the types such that types are arranged in a hierarchy. For example, every type is a subtype of `ANY`. Any atomic type is a subtype of `AnyAtomic`. Integer is a subtype of `Long`. An array is a subtype of `JSON` if its element type is `JSON` or another subtype of `JSON`.

A data item is an instance of a type (T) if the data item's type is (T) or a subtype of (T).

This relationship is important because the usual subtype-substitution rule is supported by SQL for Oracle NoSQL Database. If an operation expects input items of type (T), or produces items of type (T), then it can also operate on or produce items of type (S) if (S) is a subtype of (T). (There is an exception to this rule. See [Subtype-Substitution Rule Exceptions](#)).

The following figure illustrates this data hierarchy. Dotted boxes in the figure represent collections of types.



In addition to the subtype relationships described here, the following relationships are also defined:

- Every type is a subtype of itself. A type (T) is a proper subtype of another type (S) if (T) is a subtype of (S) and (T) is not equal to (S).
- An enum type is a subtype of another enum type if both types contain the same tokens and in the same order, in which case the types are actually considered equal.
- Timestamp(p1) is a subtype of Timestamp(p2) if $p1 \leq p2$.
- A record type (S) is a subtype of another record type (T) if:
 1. both types contain the same field names in the same order; and

2. for each field, its value type in (S) is a subtype of its value type in (T); and
 3. nullable fields in (S), are also nullable in (T).
- (Array(S)) is a subtype of (Array(T)) if (S) is a subtype of (T).
 - (Map(S)) is a subtype of (Map(T)) if (S) is a subtype of (T).

Subtype-Substitution Rule Exceptions

Ordinarily, if an operation expects input items of type (T), or produces items of type (T), then it can also operate on or produce items of type (S) if (S) is a subtype of (T). However, there are two exceptions to this rule.

The first exception concerns numeric values. Double and Float are subtypes of Number. However, Double and Float include three special values: NaN (not a number), positive infinity, and negative infinity. These values are not in the domain of Number. Therefore, an operation that expects a Number value will also work with Double/Float values as long as these values are not one of the three special values. If one of the three special values are used with Number, an error is raised.

Secondly, items whose type is a proper subtype of Array(JSON) or Map(JSON) cannot be used as:

- record/map field values if the field type is JSON, Array(JSON) or Map(JSON); or
- elements of arrays whose element type is JSON, Array(JSON) or Map(JSON).

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

For example, consider a JSON document which is a map value whose associated type is Map(JSON). The document may contain an array whose values contain only integers. However, the type associated with the array cannot be Array(integer), it must be Array(JSON). If the array had type Array(integer), the user would not be able to add any non-integer values to it.

SQL for Oracle NoSQL Database Sequences

A *sequence* is an important concept in SQL for Oracle NoSQL Database. It is used wherever expressions and operators are discussed.

A sequence is the result of any expression that returns zero or more items. Sequences are not containers; they cannot be nested.

Note that an array is not a sequence; rather it is a single item that contains other items in it.

Sequences have a type. A sequence type specifies the type of items that may appear in a sequence. For example, a sequence could contain atomic elements that are of type integer. In this case, the sequence type would be integer.

Sequence types have a cardinality. The cardinality indicates constraints on how many items can or must appear in the sequence:

- *
indicates a sequence of zero or more items.
- +
indicates a sequence of one or more items.

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

When we say that the result of an expression must have a sequence of a certain type, what we mean is the sequence must have that type or any subtype of that type.

Sequence subtypes are defined as follows:

- The empty sequence is a subtype of all sequence types whose quantifier is * or ?.
- A sequence type (s1) is a subtype of another sequence type (s2) if:
 - s1's item type is a subtype of s2's item type; and
 - s1's sequence quantifier is a sub-quantifier for s2's quantifier.

The following table shows the subtype relationship for the various quantifiers:

S2 \ S1	one	?	+	*
one	TRUE	FALSE	FALSE	FALSE
?	TRUE	TRUE	FALSE	FALSE
+	TRUE	FALSE	TRUE	FALSE
*	TRUE	TRUE	TRUE	TRUE

Sequence Concatenation Function

Use the `any* seq_concat(<argument>*)` to concatenate one sequence to another sequence. This function evaluates its arguments (if any) in the order they are listed in the arguments list, and concatenates the sequences returned by these arguments.

For an example of using this function, see [Adding Elements to an Array](#).

8

SQL for Oracle NoSQL Database Queries

This chapter describes the Select-From-Where (SFW) expression, which is the core expression used to form SQL queries. For examples of using SFW expressions, see these extended examples:

- [Simple SELECT Queries](#)
- [Working with complex data](#)
- [Working With Indexes](#)
- [Working with JSON](#)

For a more detailed description of the language see the *SQL for Oracle NoSQL Database Specification*.

Note:

The examples shown in this chapter rely on the sample data shown in [Example Data](#).

Select-From-Where (SFW) Expressions

A query is always a single Select-From-Where (SFW) expression. The SFW expression is essentially a simplified version of the SQL Select-From-Where query block. The two most important simplifications are the lack of support for joins and for subqueries. On the other hand, to manipulate complex data (records, arrays, and maps), SQL for Oracle NoSQL Database provides extensions to traditional SQL through novel kinds of expressions, such as path expressions.

The semantics of the SFW expression are similar to those in standard SQL. Processing starts with the FROM clause, followed by the WHERE clause (if any), followed by the ORDER BY clause (if any), followed by the OFFSET and LIMIT clauses, and finishing with the SELECT clause. Each clause is described below. A query must contain exactly one SFW expression, which is also the top-level expression of the query. Subqueries are not supported yet.

```
SELECT <expression>
FROM <table name>
[WHERE <expression>]
[ORDER BY <expression> [<sort order>]]
[OFFSET <number>]
[LIMIT <number>];
```

Each of the SFW clauses are introduced in the following sections. For details on each clause, see the *SQL for Oracle NoSQL Database Specification*.

SELECT Clause

SELECT clauses come in two forms. In the first form, it contains only a single star (*) symbol. This form simply returns all rows.

```
SELECT * FROM Users;
```

In the second form, the SELECT clause contains a comma-separated list of field expressions, where each expression is optionally associated with a name. In the simplest case, each expression is simply the name of a column in the table from which data is being selected.

```
SELECT id, firstname, lastname FROM Users;
```

The AS keyword can also be used:

```
SELECT id, firstname AS Name, lastname AS Surname FROM Users;
```

SELECT clauses can contain many different kinds of expressions. For more information, see [Expressions](#).

The SELECT clause always returns a record. Normally, the record has one field for each field expression, and the fields are arranged in the same order as the field expressions. Each field value is the value computed by the corresponding field expression and its name is the name associated with the field expression. If no field name is provided explicitly (using the AS keyword), one is automatically generated for you.

To create valid records, the field names must be unique, and they must return at most one item. If a field expression returns more than one result, the result is returned in an array.

If the result of a field expression is empty, NULL is used as the value of the corresponding field in the record returned by SELECT.

Note:

If the SELECT clause contains only one field expression with no associated name, then just the value returned by the clause is returned. If this value is already a record, then this is returned. If this value is not a record, then it is wrapped in a record before being returned.

SELECT Clause Hints

The SELECT clause can contain one or more hints which are used to help choose an index to use for the query. A hint is a comment that begins with a + symbol:

```
/*+ <hint> */
```

Each hint takes the form:

```
<hint type> (<table path> [<index name>]) [comment string]
```

The following hint types are supported:

- **FORCE_INDEX**
Specifies a single index, which is used without considering any of other indexes. This is true even if there are no index predicates for the forced index. However, if the query has an ORDER BY clause, and the forced index is not the sorting index, an error is thrown.
This index hint requires you to specify an <index name>.
- **PREFER_INDEXES**
The PREFER_INDEXES hint specifies one or more indexes. The query processor may or may not use one of the preferred indexes.
This index hint requires you to specify at least one <index name>.
- **FORCE_PRIMARY_INDEX**
Requires the query to use the table's primary index.
You do not specify an <index name> when you use this type of hint.
- **PREFER_PRIMARY_INDEX**
Specifies that you prefer to use the primary index for the query. This index may or may not be used.
You do not specify an <index name> when you use this type of hint.

For more information on indexes, see [Working With Indexes](#).

FROM Clause

The FROM clause is very simple: it can include only a single table. The table is specified by its name, which may be a composite (dot-separated) name in the case of child tables. The table name may be followed by a table alias.

For example, to select a table named `Users`:

```
<select expression> FROM Users <other clauses>;
```

To select a table named `People`, which is a child of a table named `Organizations`:

```
<select expression> FROM Organizations.People <other clauses>;
```

To select a table named `People` and give it the alias `u`:

```
<select expression> FROM Users u <other clauses>;
```

The result of the FROM clause is a sequence containing the rows of the referenced table. The FROM clause creates a nested scope, which exists for the rest of the SFW expression.

The SELECT, WHERE, and ORDER BY clauses operate on the rows produced by the FROM clause, processing one row at a time. The row currently being processed is called the context row. The context row can be referenced in expressions by either the table name, or the table alias.

If the table alias starts with a dollar sign (\$), then it serves as a variable declaration whose name is the alias. This variable is bound to the context row and can be referenced within the SFW expression anywhere an expression returning a single record may be used. If this variable has the same name as an external variable, it

hides the external variable. Because table alias are essentially variables, the like all other variables their names are case-sensitive.

WHERE Clause

The WHERE clause returns a subset of the rows coming from the FROM clause. Specifically, for each context row, the expression in the WHERE clause is evaluated. The result of this expression must have type BOOLEAN?. If the result is false, or empty, or NULL, the row is skipped; otherwise the row is passed on to the next clause.

For example, to limit the rows selected to just those where the column `firstname` contains `John`:

```
<select statement> <from statement> WHERE firstname = "John";
```

ORDER BY Clause

The ORDER BY clause reorders the sequence of rows it receives as input. The relative order between any two input rows is determined by evaluating, for each row, the expressions listed in the order-by clause and comparing the resulting values. Each order-by expression must have type AnyAtomic?.



Note:

It is possible to perform ordering only if there is an index that already sorts the rows in the desired order.

For detailed information on how comparison is performed for order-by expressions, see the *SQL for Oracle NoSQL Database Specification*.

For example, to order a query result by age.

```
<select statement> <from statement> WHERE firstname = "John"  
ORDER BY age;
```

It is possible to specify a sorting order: `ASC` (ascending) or `DESC` (descending). Ascending is the default sorting order. To present these results in descending order:

```
<select statement> <from statement> WHERE firstname = "John"  
ORDER BY age DESC;
```

You can also specify whether NULLS should come first or last in the sorting order. For example:

```
<select statement> <from statement> WHERE firstname = "John"  
ORDER BY age DESC NULLS FIRST;
```

Remember that ordering is only possible if there is an index that sorts the rows in the desired order. Be aware that, in the current implementation, NULLS are always sorted last in the index. The specified handling for NULLS must match the index so, currently, if the sort order is ascending then `NULL LAST` must be used, and if the sort order is descending then `NULL FIRST` must be used.

Comparison Rules

This section describes the sorting rules used when query results are sorted.

First, consider the case where only one ORDER BY clause is used in the query.

Two rows are considered equal if both rows contain the same number of elements, and for equivalent positions in each row, the atomic values are identical. So if you have two rows, R1 and R2, then they are equal if $R1[0] = R2[0]$ and $R1[1] = R1[1]$. In this context, NULLs are considered equal only to other NULLs.

Assuming that the number of elements in R1 and R2 are equal, then R1 is less than R2 if any of the following is true:

- No NULLs appear in either row and sorting is in ascending order. In this case, R1 is less than R2 if there are a positionally-equivalent pair of atomic elements (as evaluated from lowest to highest) where the R1 element is less than the R2 element. That is, if $R1[1] < R2[2]$ then R1 is less than R2.
- No NULLs appear in either row and sorting is in descending order. In this case, R1 is less than R2 if there are a positionally-equivalent pair of atomic elements (as evaluated from lowest to highest) where the R1 element is greater than the R2 element. That is, if $R1[1] > R2[2]$ then R1 is less than R2.
- A NULL appears in R2, but not in R1, and sorting is in ascending order with NULLS LAST.
- A NULL appears in R2, but not in R1, and sorting is in descending order with NULLS FIRST.

If multiple ORDER BY statements are offered, then atomic values are returned for comparison purposes by evaluating the statements from left to right.

Be aware that if an expression returns an empty sequence, then the return value is NULL.

If no sorting order is provided to the query, then by default ascending order with NULLS LAST is used. If only the sort order is specified, then NULLs sort last if the order is ascending. Otherwise, they sort first.

OFFSET Clause

Specifies the number of initial query results that should be skipped; that is, they are not returned. This clause accepts a single non-negative integer as its argument. This argument may be a single integer literal, or a single external variable, or any expression which is built from literals and external variables.

Although it is possible to use this clause without an ORDER BY clause, it does not make sense to do so. Without an ORDER BY clause, results are returned in random order, so the subset of results skipped will be different each time the query is run.

LIMIT Clause

Specifies the maximum number of results to return. This clause accepts a single non-negative integer as its argument. This argument may be a single integer literal, or a single external variable, or any expression which is built from literals and external variables.

Although it is possible to use this clause without an ORDER BY clause, it does not make sense to do so. Without an ORDER BY clause, results are returned in random order, so the subset of results returned will be different each time the query is run.

9

Expressions

In general, an expression represents a set of operations to be executed in order to produce a result. Expressions are built by combining other subexpressions using operators (arithmetic, logical, value and sequence comparisons), function calls, or other grammatical constructs. The simplest kinds of expressions are constants and references to variables or identifiers.

In SQL for Oracle NoSQL Database, the result of any expression is always a sequence of zero or more items. Notice that a single item is considered equivalent to a sequence containing that single item.

Note:

The examples shown in this chapter rely on the sample data shown in [Example Data](#).

Path Expressions

To navigate inside complex values and select their nested values, SQL for Oracle NoSQL Database supports path expressions. A path expression has an input expression followed by one or more steps.

```
<primary_expressions>.<step>*
```

Note:

A path expression over a table row must always start with the table's name or the table's alias (if one was included in the FROM clause).

There are three kinds of path expression steps: field, filter, and slice steps. Field steps are used to select field/entry values from records or maps. Filter steps are used to select array or map entries that satisfy some condition. Slice steps are used to select array entries based on their position inside the containing array. A path expression can mix different kinds of steps.

All steps iterate over their input sequence, producing zero or more items for each input item. If the input sequence is empty, the result of the step is also empty. Otherwise, the overall result of the step is the concatenation of the results produced for each input item. The input item that a step is currently operating on is called the *context item*, and it is available within the step expression using the dollar sign (\$) variable. This context-item variable exists in the scope created by the step expression.

For all steps, if the context item is NULL, it is just added into the output sequence with no further processing.

In general, path expressions may return more than one item as their result. Such multi-item results can be used as input in two other kinds of expressions: sequence-comparison operators and array constructors.

Field Step Expressions

A field step selects the value of a field from a record or map. The field to select is specified by its field name, which is either given explicitly as an identifier, or is computed by a name expression. The name expression must be of type string.

```
<primary_expression>.<id> | <string> | <var_ref> |  
<parenthesized_expr> | <func_call>*
```

As a simple example, the field step expression `u.address.city`:

```
SELECT id, u.address.city FROM Users u;
```

Retrieves the field "city" from the "address" column in the Users ("u") table.

A field step processes each context item as follows:

1. If the context item is an atomic item, it is skipped (the result is empty).
2. The name expression is evaluated. If the name expression returns the empty sequence or NULL, the context item is skipped. Otherwise, the evaluated name expression is passed to the next step.
3. If the context item:
 - Is a record
and if that record contains a field identical to the evaluated name expression, then that field is returned. Otherwise, an error is raised.
 - Is a map
and if that map contains a field identical to the evaluated name expression, then that field is returned. Otherwise, an empty result is returned.

If the context item (\$) is an array, then the field step is applied to each element of the array with the context item being set to the current array element. If the context item is an atomic item, it is skipped (the result is empty).

Map Filter Step Expressions

A map filter step is used with records and maps to select either the field name (keys) or the field values of the fields that satisfy a given condition. This condition is specified as a predicate expression inside parentheses. If the predicate expression is missing, it is assumed to be `true` — all the field names or values are returned.

```
<primary_expression>.keys | values (<predicate>)
```

where `keys` references the record's or map's field name, and `values` references the record's or map's field values.

In addition to the context-item variable (\$), the predicate expression may reference the following two variables: `$key` is bound to the name of the context field — that is, the current field in \$, and `$value` is bound to the value of the context field. The predicate expression must be boolean.

A simple example is `u.expenses.keys($value > 1000)`, which selects all the expenses greater than \$1000. Combined with this query:

```
SELECT id, u.expenses.keys($value > 1000) FROM Users u;
```

all the user IDs and expense fields are returned where more than 1000 was spent.

A map filter step processes each context item as follows:

1. If the context item is an atomic item, it is skipped (the result is empty).
2. If the context item is a record or map, the step iterates over its fields. For each field, the predicate expression is evaluated. A NULL or an empty result from the predicate expression is treated as a false value. If the predicate result is true, the context field is selected and either its name or its value is returned; otherwise the context field is skipped.

 **Note:**

If the context item (\$) is an array, then the map filter step is applied to each element of the array with the context item being set to the current array element.

Array Filter Step Expressions

An array filter step is used with arrays to select elements of arrays by evaluating a predicate expression for each element. Elements are selected or rejected depending on the results of the predicate expression. If the predicate expression is missing, it is assumed to be `true` — all the array elements are returned.

```
[<primary_expression>[<predicate_expression>]]
```

Notice in the syntax that the entire expression is enclosed in square brackets (`[]`). This is the array constructor. Use of the array constructor is frequently required in order to obtain the desired result, and so we show it here. The use of the explicit array constructor guarantees that the records in the result set will always have an array as their second field. For example:

```
SELECT lastName,  
[ u.address.phones[$element.area = 650].number ] AS phoneNumbers  
FROM Users u;
```

Assume that `u.address.phones` references one or more phone numbers. Without the array constructor, the result records would contain an array for users with more than one phone (because the information would be held in an array in the store anyway), but just a single integer for users with just one phone. For users with just one phone, the `phones` field might not be an array (containing a single phone object), but just a single phone object. If such a single phone object has area code 650, its number will be selected, as expected.

In addition to the context-item variable (`$`), the predicate expression may reference the following two variables: `$element` is bound to the current element in `$`, and `$pos` is bound to the position of the context element within `$`. Positions are counted starting with 0.

An array filter step processes each context item as follows:

1. If the context item is not an array, an array is created and the context item is added to that array. Then the array filter is applied to this single-item array.
2. If the context item is an array, the step iterates over the array elements and computes the predicate expression on each element.

The predicate expression must return a boolean item, or a numeric item, or the empty sequence, or NULL. A NULL or an empty result from the predicate expression is treated as a false value. If the predicate result is true/false, the context element is selected/skipped, respectively. If the predicate result is a number, the context element is selected only if \$pos equals that number. This means that if the predicate result is a negative number, or greater or equal to the array size, the context element is skipped.

Array Slice Step Expressions

An array slice step is used with arrays to select elements of arrays based on element position. The elements to select are identified by specifying boundary positions which identify "low" position and "high" positions. Each boundary expression must return at most one item of type LONG or INTEGER, or NULL. The low and/or the high expression may be missing. The context-item variable (\$) is available during the computation of the boundary expressions.

```
<primary_expression>[<low>:<high>]
```

For example, assume an array of connects ordered from the strongest connect (position 0) to the weakest, select the strongest connection for the user with id 10:

```
select connections[0] as strongestConnection from Users
where id = 10;
```

Select user 10's five strongest connections, and return the array (notice the use of the array constructor):

```
select [ connections[0:4] ] as strongConnections from Users
where id = 10;
```

Select user 10's five weakest connections:

```
select [ connections[size($) - 5 : ] ] as weakConnections from Users
where id = 10;
```

An array slice step processes each context item as follows:

1. If the context item is not an array, an array is created and the context item is added to that array. Then the array filter is applied to this single-item array.
2. If the context item is an array, the boundary expressions (if any) are evaluated.

If any boundary expression returns NULL or an empty result, the context item is skipped.

Otherwise, if the low expression is absent, or if it evaluates to less than 0, the lower boundary is set to 0. If the high expression is absent, or if it evaluates to higher than the array_size - 1, it is set to array_size - 1.

3. After the low and high positions are determined, the step selects all the elements positions, inclusively, between those two boundaries. If the low position is greater than the high position, then no elements are selected.

Constant Expressions

There are five kinds of constants available:

- **Strings.**
Sequences of unicode characters enclosed in double or single quotes. String literals are translated into String items.
- **Integer numbers**
Sequences of one or more digits. Integer literals are translated into Integer items, if their value fits in 4 bytes, otherwise into Long items.
- **Real numbers**
Representation of real numbers using “dot notation” and/or exponent. Real literals are translated into Double items.
- The boolean values `true` and `false`.
- The JSON null value.

Column Reference Expression

A column-reference expression returns the item stored in the specified column within the context row (the row that a WHERE, ORDER BY, or SELECT clause is currently working on).

A column-reference expression consists of one identifier, or two identifiers separated by a dot. If there are two ids, the first is considered to be a table name or /alias, and the second a column in that table. A single id refers to a column in the table referenced inside the FROM clause.

Notice that child tables in Oracle NoSQL Database have composite names using dot as a separator among multiple ids. As a result, a child-table name cannot be used in a column-reference expression; instead, a table alias must be used to access a child table column using the two-id format. For example, if "Address" is a child table of Persons, then:

```
SELECT id, p.Address.state FROM Persons p;
```

Variable Reference Expression

A variable reference expression is: `$(variablename)`.

A variable-reference expression returns the item that the specified variable is currently bound to. Syntactically, a variable-reference expression is just the name of the variable.

Examples of variable usage can be found scattered throughout [Working with complex data](#). For example, [Working With Maps](#) discusses and shows the usage of several implicitly declared variables related to maps (`$key` and `$value`).

Searched Case Expressions

```
CASE
  WHEN <expr> THEN <expr>
  (WHEN <expr> THEN <expr>)*
  (ELSE <expr>)?
END;
```

The searched case expression consists of a number of when-then pairs, followed by an optional `else` clause at the end. Each when expression is a condition that must return boolean. The `then` expressions as well as the `else` expression may return any sequence of items.

The case expression is evaluated by:

1. Evaluating the `when` expressions from top to bottom until the first one is discovered that returns true.
2. The `then` expression for the previously identified `when` is evaluated. This result is returned as the result for the entire case expression.
3. If no `when` expression returns true, but there is an `else` expression, then that expression is evaluated and its result is the result of the entire case expression.
4. Otherwise, the result of the entire case expression is the empty sequence.

For example, construct a map using the map constructor (`{}`) in which the `phones:` element is either the contents of the `phones` column, or a string to indicate nothing was found in that column:

```
select {
  "last_name" : u.lastName,
  "phones" : case
    when exists u.address.phones then u.address.phones
    else "Phone info absent at the expected place"
  end,
  "high_expenses" : [ u.expenses.keys($value > 5000) ]
}
from Users u;
```

For more examples of using searched case expressions, see [Using Searched Case](#).

Cast Expressions

The cast expression creates, if possible, new items of a given target type from the items of its input sequence.

```
CAST (<input_sequence> AS <target_type><quantifier>)
```

Cast expressions are evaluated as follows:

1. A cardinality check is performed based on the `<quantifier>`. If `<quantifier>` is:
 - `*`
then `<input_sequence>` may have any number of items.
 - `+`
then `<input_sequence>` must have at least one item.

- ?
then <input_sequence> must have at most one item.
- No quantifier
then <input_sequence> must exactly one item.

If this check fails, an error is raised.

2. Each input item is cast to the <target_type> according to the following recursive rules:
 - If the type of the input item is equal to the target item type, the cast is a noop: the input item itself is returned.
 - If the target type is a wildcard type, the cast is a noop if the type of the input item is a subtype of the wildcard type; otherwise an error is raised.
 - If the target type is JSON, then:
 - an error is raised if the input item has a non-json atomic type.
 - if the input item has a type that is a json atomic type or ARRAY(JSON) or MAP(JSON), the cast is a noop.
 - if the input item is a non-json array, a new array of type ARRAY(JSON) is constructed, each element of the input array is cast to JSON, and the resulting item is appended into the new json array.
 - if the input item is a non-json map, a new map of type MAP(JSON) is constructed, each field value of the input map is cast to JSON, and the resulting value together with the associated field name are inserted into the new json map.
 - if the input item is a record, it is cast to a map of type MAP(JSON).
 - If the target type is an array type, an error is raised if the input item is not an array. Otherwise, a new array is created whose type is the target type. Each element in the input array is cast to the element type of the target array, and each such item is appended into the new array.
 - If the target type is a map type, an error is raised if the input item is not a map or a record. Otherwise, a new map is created whose type is the target type. Each element in the input map/record is cast to the element type of the target map, and the resulting value together with the associated field name is inserted into the new map.
 - If the target type is a record type, an error is raised if the input item is not a record or a map. Otherwise, a new record is created whose type is the new target type.

If the input item is a record, its type must have the same fields and in the same order as the target type. In this case, each field value in the input record is cast to the value type of the corresponding field in the target type, and the resulting field value together with the associated field name is added to the new record.

If the input item is a map, then for each map field, if the field name exists in the target type, the associated field value is cast to the value type of the corresponding field in the target type, and the resulting field value together with the associated field name is added to the new record. Any fields in the new record whose names do not appear in the input map have their associated field values set to their default values.

- If the target type is a string, the input item may be of any type. In other words, every data type can be cast to a string. For complex items their “string value” is a json-text representation of their value. For timestamps, their string value is in UTC and has the format `yyyy-MM-dd['T' HH:mm:ss]`. For binary items, their string value is a base64 encoding of their bytes.
- If the target type is an atomic type other than string, the input item must also be atomic. Among atomic items and types the following casts are allowed:
 - Every numeric item can be cast to every other numeric type. The cast is done as in Java.
 - String items may be cast-able to all other atomic types. Whether the cast succeeds 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 cast-able to all the timestamp types. If the target type has a smaller precision than the input item, the resulting timestamp is the one closest to the input timestamp in the target precision. For example, consider the following 2 timestamps with precision 3: `2016-11-01T10:00:00.236` and `2016-11-01T10:00:00.267`. The result of casting these timestamps to precision 1 is: `2016-11-01T10:00:00.2` and `2016-11-01T10:00:00.3`, respectively.

For examples of using the cast expression, see [Casting Datatypes](#).

10

Operators

This chapter describes the various operators you can use with your SQL expressions.

Note:

The examples shown in this chapter rely on the sample data shown in [Example Data](#).

Logical Operators

The binary AND and OR operators and the unary NOT operator have the usual semantics. Their operands are conditional expressions, which must have type BOOLEAN.

An empty result from an operand is treated as the false value. If an operand returns NULL then:

- The AND operator returns false if the other operand returns false; otherwise, it returns NULL.
- The OR operator returns false if the other operand returns false; otherwise, it returns NULL.
- The NOT operator returns NULL.

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. The available value comparison operators are:

- =
- !=
- >
- >=
- <
- <=

If any operand returns more than one item, an error is raised. If both operands return the empty sequence, the operands are considered equal (so true will be returned if the operator is =, <=, or >=). If only one of the operands returns empty, the result of the comparison is false unless the operator is !=.

Among atomic items, if the types of the items are not comparable, false is returned. The following rules defined what atomic types are comparable and how the comparison is done in each case.

- A numeric item is comparable with any other numeric item. If an integer/long value is compared to a float/double value, the integer/long will first be cast to float/double.
- A string item is comparable to another string item.
A string item is also comparable to an enum item. In this case, before the comparison the string is cast to an enum item in the type of the other enum item. Such a cast is possible only if the enum type contains a token whose string value is equal to the source string. If the cast is successful, the two enum items are then compared as explained in the next bullet; otherwise, the two items are incomparable and false is returned.
- Two enum items are comparable only if they belong to the same type. If so, the comparison is done on the ordinal numbers of the two enums (not their string values).
- Binary and fixed binary items are comparable with each other for equality only. The 2 values are equal if their byte sequences have the same length and are equal byte-per-byte.
- A boolean item is comparable with another boolean item.
- A timestamp item is comparable to another timestamp item, even if their precisions are different.
- JNULL (JSON null) is comparable with JNULL. If the comparison operator is !=, JNULL is also comparable with every other kind of item, and the result of such a comparison is always true, except when the other item is also JNULL.

The semantics of comparisons among complex items is:

- A record is comparable with another record for equality only, and only if they contain comparable values. To be equal, the 2 records must have equal sizes (number of fields) and for each field in the first record, there must exist a field in the other record such that the two fields are at the same position within their containing records, have equal field names, and equal values.
- A map is comparable with another map for equality only, and only if they contain comparable values. To be equal, the 2 maps must have equal sizes (number of fields) and for each field in the first map, there must exist a field in the other map such that the two fields have equal names and equal values.
- An array is comparable to another array, if the elements of the 2 arrays are comparable pair-wise. Comparison between 2 arrays is done lexicographically. That is, the arrays are compared like strings, with the array elements playing the role of the "characters" to compare.

As with atomic items, if two complex items are not comparable according to the above rules, false is returned. Comparisons between atomic and complex items return false always.

 **Note:**

The reason for returning false for incomparable items, instead of raising an error, is to handle schema-less applications where different table rows may contain very different data, or differently shaped data. As a result, even the writer of the query may not know what kind of items an operand may return and an operand may indeed return different kinds of items from different rows. Nevertheless, when the query writer compares “something” with, say, an integer, they expect that the “something” will be an integer and they would like to see results from the table rows that fulfill that expectation, instead of the whole query being rejected because some rows do not fulfill the expectation.

Sequence Comparison Operators

Comparisons between two sequences is done using the following operators:

- =any
- !=any
- >any
- >=any;
- <any
- <=any

The result of an any operator on two input sequences S1 and S2 is true only if:

1. There is a pair of items, i1 and i2; and
2. i1 belongs to S1, and i2 belongs to S2; and
3. i1 and i2 compare true using the corresponding value-comparison operator.

Otherwise, if any of the input sequences contains NULL, the result is NULL.

Otherwise, the result is false.

IS NULL Operator

The IS NULL operator test whether the result of its input expression is SQL NULL. (SQL NULL is used when a non-JSON field is set to NULL.) IS NULL requires an input expression that returns a single item. If that single item is SQL NULL, then IS NULL returns true.

A table field can be NULL if it is explicitly set to NULL, or if the table field is simply not populated when you import data.

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. This means that IS NULL cannot be used to identify missing fields from JSON columns. Use the [Exists Operator](#) instead.

 **Note:**

IS NULL returns false for JSON fields which exist but are set to NULL. This is because for JSON data, NULL is the JSON NULL, not the SQL NULL.

IS NOT NULL can also be used. It is equivalent to:

```
NOT (IS NULL <expression>)
```

For an example of using IS NULL and IS NOT NULL, see [Filtering results](#).

Exists Operator

The exists operator checks whether a sequence is empty. True is returned if the sequence is not empty.

Note that this operator returns true even if the sequence contains null data. So for strongly typed data (that is, non-JSON data columns), this operator will always return true because on data import the column will be always at minimum populated with the SQL NULL. To check whether strongly typed data is NULL, use the [IS NULL Operator](#).

For an examples of using the exists operator, see [Using WHERE EXISTS with JSON](#).

Is-Of-Type Operator

Returns true if the input sequence matches the identified type.

```
<sequence> IS OF TYPE (<type>*|+|?,<type>*|+|?,...)
```

or

```
<sequence> IS NOT OF TYPE (<type>*|+|?,<type>*|+|?,...)
```

The is-of-type operator checks the input sequence type against one or more target sequence types. It returns true if both of the following conditions are true:

- The cardinality of the input sequence matches the quantifier of the target type:
 - If the quantifier is *, the input sequence may have any number of items.
 - If the quantifier is +, the input sequence must have at least one item.
 - If the quantifier is ?, the input sequence must have at most one item.
 - If there is no quantifier, the input sequence must have exactly one item.
- All of the items in the input sequence are instance of the specified type(s). For the purpose of this check, a NULL is not considered to be an instance of any type.

SQL for Oracle NoSQL Database's subtype-supertype relationship model is relevant to the usage of this operator. See [Type Hierarchy](#) for more information.

If the cardinality requirement is met and the input sequence contains a NULL, this operator returns NULL. In all other cases, the result is false.

 **Note:**

If the number of the target types is greater than one, the expression is equivalent to OR-ing that number of is-of-type expressions, each having one target type.

For an example of using is-of-type, see [Examining Data Types JSON Columns](#).

11

Constructors

SQL for Oracle NoSQL Database offers two constructors that you can use: Array and Map constructors.

Note:

The examples shown in this chapter rely on the sample data shown in [Example Data](#).

Array Constructors

```
[ <expression>, <expression>, ... ]
```

An array constructor constructs a new array out of the items returned by the expressions inside the square brackets. These expressions are computed left to right and the produced items are appended to the array. Any NULLs produced by the input expressions are skipped (arrays cannot contain NULLs).

The type of the constructed array is determined during query compilation, based on the types of the input expressions and the usage of the constructor expression. Specifically, if a constructed array can be inserted in another constructed array and this “parent” array has type ARRAY(JSON), then the “child” array will also have type ARRAY(JSON). This is because “typed” data is not allowed inside JSON data.

For example, the use of the explicit array constructor here means that the field will exist in all returned rows, even if the path inside the array constructor returns empty. Without this constructor, a NULL can be returned for the field. With it, an empty result returns an empty array.

```
select firstname, lastname,  
[u.expenses.keys($value > 1000)] AS Expenses  
from Users u;
```

Map Constructors

```
{ <expression>:<expression>,  
<expression>:<expression>, ... }
```

A map constructor constructs a new map out of the items returned by the expressions inside the curly brackets. These expressions come in pairs: each pair computes one field.

The first expression in a pair must return at most one string, which serves as the field's name. If the field name expression returns the empty sequence, no field is constructed.

The second expression returns the field value. If this expression returns more than one item, an array is implicitly constructed to store the items, and that array becomes the field value. If the field value expression returns the empty sequence, no field is constructed.

If the computed name or value for a field is NULL the field is skipped (maps cannot contain NULLs).

The type of the constructed map is determined during query compilation, based on the types of the input expressions and the usage of the constructor expression. Specifically, if a constructed map can be inserted in another constructed map and this "parent" map has type MAP(JSON), then the "child" map will also have type MAP(JSON). This is because "typed" data is not allowed inside JSON data.

For example, construct a map consisting of user's last name and their phone numbers:

```
select {
  "last_name" : u.lastName,
  "phones" : u.address.phones
}
from Users u;
```

12

Built-in Functions

You can use function-call expressions to invoke functions, which currently can only be built-in (system) functions. The available built-in system functions are as follows:

- Size function.
Returns the size (number of fields/entries) of a complex item such as a record, array, or map. For an example of usage, see [Using the size\(\) Function](#)
- Timestamp functions.
Returns specific information from a Timestamp data type, such as the year or month that the data represents. See [Timestamp Functions](#) for a description of those functions. See [Working with Timestamps](#) for an example of how to use the Timestamp functions.
- Sequence concatenation function.
Used to concatenate one sequence to another. See [Sequence Concatenation Function](#) for a description of this function. See [Adding Elements to an Array](#) for an example of how to use this function.
- Time to Live functions.
Extracts Time to Live properties for table rows. See [Time to Live Functions](#) for a list of these functions.
- Time functions.
Miscellaneous functions related to the system clock. See [Time Functions](#) for a list of these functions.

Time to Live Functions

A time to live (TTL) value can be set on a table by table, or row by row basis. Data that reaches its time to live value is said to have *expired*, and it will no longer be returned as the result of a query. For information on setting TTL values on a row by row basis, see [Managing Time to Live Values](#).

TTL properties are not a normal part of table schema. They are not stored as top-level columns or nested fields, and as such cannot be queried using ordinary query mechanisms. Instead, you can use the following functions to examine a row's TTL values:

- `<integer> remaining_hours(<row>)`
Returns the number of full hours remaining until the row expires. A negative number is returned if the row has no expiration time.
- `<integer> remaining_days(<row>)`
Returns the number of full days remaining until the row expires. A negative number is returned if the row has no expiration time.
- `<timestamp(0)> expiration_time(<row>)`

Returns the expiration time of the row as a timestamp value of precision zero. If the row has no expiration time, this function returns `1970-01-01T00:00:00.000`.

- `<long> expiration_time_millis(<row>)`

Returns the expiration time of the row as the number of milliseconds since January 1, 1970 UTC. Zero (0) is returned if the row has no expiration time.

All of these functions require a row as input. The only expression that returns a row is a row variable; that is, a table alias whose name starts with `$`.

For an example of using these functions, see [Managing Time to Live Values](#).

Time Functions

The following functions can be used to return time information from the system clock:

- `<long> current_time_millis()`

Returns the current time in UTC as the number of milliseconds since January 1, 1970 UTC.

- `<timestamp(3)> current_time()`

Returns the current time in UTC as a timestamp value with millisecond precision.

For an example of using `current_time()`, see [Updating Existing Map Elements](#), and [Managing Time to Live Values](#).

13

SQL UPDATE Statements

As with standard SQL, you can use SQL for Oracle NoSQL Database to update a table row. However, SQL for Oracle NoSQL Database includes extensions to handle the richer data model that Oracle NoSQL Database offers.

When you execute an UPDATE statement using SQL for Oracle NoSQL Database, the update takes place directly, resulting in a very efficient way to update data. Using UPDATE, you do not need to retrieve data from the store, modify it, and then write the data back to the store. Instead, you send the UPDATE statement to the store, which directly updates the row without any further network traffic.

For examples of using update statements, see [Modifying Table Rows using UPDATE Statements](#).

Update Statement Syntax

The update statement syntax is:

```
[<prolog>]
UPDATE <table_name> [AS <table_alias>]
    <update_clause>[, <update_clause>]*
WHERE <expr>
[<returning_clause>];
```

where:

- <prolog> is optional, and can be used to declare external variables.
- <table_name> is the name of the table you are updating.
- <table_alias> is optional, but it can be omitted only if top-level columns are being accessed by the statement. Otherwise, as is the case for read-only queries, an alias is required as the first step in path expressions that access nested fields.
- <update_clause> describe one or more update actions to take on the table. These are described in detail in the remainder of this chapter.
- The WHERE clause specifies which row to update. Only single row updates are allowed, so the WHERE clause must specify an exact primary key.
- <returning_clause> is optional, and it indicates what should be returned as a result of the update. If it is not specified, the number of rows updated is returned. If no rows match the WHERE clause, then 0 is returned. If a match is found, then 1 is always returned because update statements can only update a single row at a time.

The <returning_clause> can also act as a SELECT clause. If * is specified, then the entire updated row is returned. Otherwise, the <returning_clause> can provide a list of expressions that indicate what should be returned from the updated table.

The rest of this chapter describes the <update_clause> in detail.

Update Clauses

Update clauses are used to describe what modifications are to be made to a table row. There are five types of clauses that you can use:

- **SET**
Updates the value of one or more fields. See [SET Clause](#) for details.
- **ADD**
Adds new elements to one or more arrays. See [ADD Clause](#) for details.
- **PUT**
Adds new fields to one or more maps, or updates the value of an existing map field. See [PUT Clause](#) for details.
- **REMOVE**
Removes elements/fields from one or more arrays/maps. See [REMOVE Clause](#) for details.
- **SET TTL**
Used to modify the row's expiration time. See [SET TTL Clause](#) for details.

Multiple clauses and a combination of clauses can be used by separating it by comma. For example:

```
update_clause :
(SET set_clause (COMMA (update_clause | set_clause))* |
(ADD add_clause (COMMA (update_clause | add_clause))* |
(PUT put_clause (COMMA (update_clause | put_clause))* |
(REMOVE remove_clause (COMMA remove_clause))* ) ;
```

SET Clause

```
set <target_expr> = <new-value_expr>
```

The SET clause changes the value of existing information in the targeted table. Its target expression which identifies the information to change. Its new-value expression identifies what the targeted information will become.

- **Target Expression**
A target expression can be atomic or complex, but it will always be nested inside a complex item (its parent item). For each such target expression, the new-value expression is evaluated and the new-value results are used to replace the target item.

If the target expression returns a NULL item, then either the target item itself is the NULL item or one of its ancestors is NULL. If the target item is NULL, the result of the new-value expression are used to replace the NULL. If one of the target item's ancestors are NULL, then the entire clause is a noop.
- **New-Value Expression**
The new-value expression can return zero or more items. If it returns an empty result, the SET is a noop. If it returns more than one item, the items are enclosed inside a newly constructed array in the same way as the way the SELECT clause treats multi-valued expressions in the select list. The result of the new-value

expression is then cast to the type expected by the parent item for the target field. (See [Cast Expressions](#) for details.) If the cast fails, an error is raised; otherwise the new item replaces the target item within the parent item.

The new-value expression may reference the implicitly declared variable \$, which is bound to the current target item. Use of the \$ variable makes it possible to have target expressions that return more than one item. In this case the SET clause will iterate over the set of target items, bind the \$ variable for each target item, compute the new-value expression, and replace the target item with the result of the new-value expression.

For examples of using a SET clause, see [Changing Field Values](#), [Changing an Existing Element in an Array](#), and [Updating Existing Map Elements](#)

ADD Clause

```
add <target_expr> <position_expression> <new-elements_expr>
```

or

```
add <target_expr> <new-elements_expr>
```

The ADD clause is used to add new elements into one or more arrays. It consists of:

- a target expression, which should return one or more array items,
- an optional position expression, which specifies the position within the array where the new element(s) should be placed at,
- and a new-elements expression that returns the new elements to insert.

This clause iterates over the sequence returned by the target expression. For each target item, if the item is not an array it is skipped. Otherwise, the position expression (if present) and the new-elements expression are computed for the current target array. These two expressions may reference the \$ variable, which is bound to the current target array.

If the position expression is missing, or if it returns an empty result, the new elements are appended at the end of the target array. An error is raised if the position expression returns more than one item or a non-numeric item. Otherwise, the returned item is cast to an integer. If this integer is less than 0, it is set to 0. If it is greater or equal to the array size, the new elements are appended.

If the the new-values expression returns nothing, the ADD clause is a noop. Otherwise, each item returned by this expression is cast to element type of the array. An error is raised if any of these casts fails. Otherwise, the new elements are inserted into the target array at the indicated position.

For examples of using an ADD clause, see [Adding Elements to an Array](#).

PUT Clause

```
put <target_expression> <new-fields_expression>
```

The PUT clause is used to add new fields into one or more maps. It consists of:

- a target expression which should return one or more map items,
- and a new-fields expression that returns one or more maps or records. These are the fields that are inserted in the target maps.

The PUT clause iterates over the sequence returned by the target expression. For each target item, if the item is not a map it is skipped. Otherwise, the new-fields expression is computed for the current target map. The new-maps expression may reference the \$ variable, which is bound to the current target map.

If the the new-fields expression returns nothing, the PUT is a noop. Otherwise, for each item returned by the new-fields expression, if the item is not a map or a record, it is skipped. If the item is a map or record, the fields of the map/record are “merged” into the current target map. This merge operation will insert a new field into the target map if the target map does not already have a field with the same key. Otherwise it will set the value of the target field to the value of the new field.

For examples of using a PUT clause, see [Adding Elements to a Map](#).

REMOVE Clause

```
remove <target_expression>
```

The remove clause consists of a single target expression, which computes the items to be removed. The REMOVE clause iterates over the target items, and for each item:

- If its parent is a record, an error is raised.
- If the target item is not NULL, it is removed from its parent.
- If the target item is NULL, it is skipped.

Note that if the target item is NULL, then one of its ancestors must be NULL. This is because arrays and maps cannot contain NULLs. Consequently, the target item is skipped because of the NULL ancestor.

For examples of using the REMOVE clause, see [Removing Elements from Arrays](#) and [Removing Elements from a Map](#).

SET TTL Clause

```
set TTL <add_expression> HOURS|DAYS
```

or

```
set TTL USING TABLE DEFAULT
```

Every table row has an expiration time that is specified in terms of a Time to Live (TTL) value. TTL values are specified as a number of days or hours. If zero (0), the row will never expire. If a row expires, its data can no longer appear in query results.

The expiration time for a row is always computed when the row is first inserted into a table. However, you can use the SET TTL clause to specify a new expiration time. The value you specify for this clause is used to compute the new expiration time.

The SET TTL clause comes in two flavors. The first contains an expression which computes a new TTL value as follows:

1. If the result of this expression is empty, the SET TTL clause is a noop.
2. If the expression result is not empty, it is cast to an integer.
3. If the resulting integer is negative, it is set to 0. This means the table row will never expire.

4. If the resulting integer is non-negative, it must be followed by a unit designation of either `HOURS` or `DAYS`.
5. The new expiration time is computed based on the current time (in UTC) plus the number of hours/days computed in `<add_expression>` rounded up to the next full hour/day. That is, if the current time is `2017-06-01T10:05:30.0` and the TTL value evaluates to 3 hours, the expiration time will be `2017-06-01T14:00:00.0`.

For an example of using the SET TTL clause, see [Managing Time to Live Values](#).

A

Introduction to the SQL for Oracle NoSQL Database Shell

This appendix describes how to configure, start and use the SQL for Oracle NoSQL Database Shell to execute SQL statements. Then, the available shell commands are described.

You can use the shell to directly execute DDL, DML, user management, security, and informational statements.

Running the shell

The shell is run interactively or used to run single commands. The general usage to start the shell is:

```
java -jar KVHOME/lib/sql.jar
      -helper-hosts <host:port[,host:port]*> -store <storeName>
      [-username <user>] [-security <security-file-path>]
      [-timeout <timeout ms>]
      [-consistency <NONE_REQUIRED(default) |
      ABSOLUTE | NONE_REQUIRED_NO_MASTER>]
      [-durability <COMMIT_SYNC(default) |
      COMMIT_NO_SYNC | COMMIT_WRITE_NO_SYNC>]
      [single command and arguments]
```

where:

- `-consistency`
Configures the read consistency used for this session.
- `-durability`
Configures the write durability used for this session.
- `-helper-hosts`
Specifies a comma-separated list of hosts and ports.
- `-store`
Specifies the name of the store.
- `-timeout`
Configures the request timeout used for this session.
- `-username`
Specifies the username to login as.

For example, you can start the shell like this:

```
java -jar KVHOME/lib/sql.jar
-helper-hosts node01:5000 -store kvstore
sql->
```

The above command assumes that a store "kvstore" is running at port 5000. You can now execute queries. In the next part of the book, you will find an introduction to SQL for Oracle NoSQL Database and how to create these query statements.

If you want to import records from a file in either JSON or CSV format, you can use the import command. For more information see [import](#).

If you want to run a script file, you can use the "load" command. For more information see [load](#).

For a complete list of the utility commands accessed through "java -jar" <kvhome>/lib/sql.jar <command>" see [Shell Utility Commands](#)

Configuring the shell

You can also set the shell start-up arguments by modifying the configuration file `.kvclirc` found in your home directory.

Arguments can be configured in the `.kvclirc` file using the `name=value` format. This file is shared by all shells, each having its named section. `[sql]` is used for the Query shell, while `[kvcli]` is used for the Admin Command Line Interface (CLI).

For example, the `.kvclirc` file would then contain content like this:

```
[sql]
helper-hosts=node01:5000
store=kvstore
timeout=10000
consistency=NONE_REQUIRED
durability=COMMIT_NO_SYNC
username=root
security=/tmp/login_root

[kvcli]
host=node01
port=5000
store=kvstore
admin-host=node01
admin-port=5001
username=user1
security=/tmp/login_user
admin-username=root
admin-security=/tmp/login_root
timeout=10000
consistency=NONE_REQUIRED
durability=COMMIT_NO_SYNC
```

Shell Utility Commands

The following sections describe the utility commands accessed through "java -jar" <kvhome>/lib/sql.jar <command>".

The interactive prompt for the shell is:

```
sql->
```

The shell comprises a number of commands. All commands accept the following flags:

- `-help`

Displays online help for the command.

- ?

Synonymous with `-help`. Displays online help for the command.

The shell commands have the following general format:

1. All commands are structured like this:

```
sql-> command [arguments]
```

2. All arguments are specified using flags that start with "-"
3. Commands and subcommands are case-insensitive and match on partial strings(prefixes) if possible. The arguments, however, are case-sensitive.

connect

```
connect -host <hostname> -port <port> -name <storeName>
[-timeout <timeout ms>]
[-consistency <NONE_REQUIRED(default) |
                        ABSOLUTE | NONE_REQUIRED_NO_MASTER>]
[-durability <COMMIT_SYNC(default) |
             COMMIT_NO_SYNC | COMMIT_WRITE_NO_SYNC>]
[-username <user>] [-security <security-file-path>]
```

Connects to a KVStore to perform data access functions. If the instance is secured, you may need to provide login credentials.

consistency

```
consistency [[NONE_REQUIRED | NONE_REQUIRED_NO_MASTER |
             ABSOLUTE] [-time -permissible-lag <time_ms> -timeout <time_ms>]]
```

Configures the read consistency used for this session.

describe

```
(describe | desc) as json
  table <table_name> (<field_name> (,<field_name>)*)?
  | index <index_name> on <table_name>
```

Provides a JSON description of a table or index.

durability

```
durability [[COMMIT_WRITE_NO_SYNC | COMMIT_SYNC |
             COMMIT_NO_SYNC] | [-master-sync <sync-policy> -replica-sync <sync-policy>
 -replica-ask <ack-policy>]] <sync-policy>: SYNC, NO_SYNC, WRITE_NO_SYNC
<ack-policy>: ALL, NONE, SIMPLE_MAJORITY
```

Configures the write durability used for this session.

exit

```
exit | quit
```

Exits the interactive command shell.

help

```
help [command]
```

Displays help message for all shell commands and sql command.

history

```
history [-last <n>] [-from <n>] [-to <n>]
```

Displays command history. By default all history is displayed. Optional flags are used to choose ranges for display.

import

```
import -table <name> -file <name> [JSON | CSV]
```

Imports records from the specified file into the named table. The records can be in either JSON or CSV format. If the format is not specified JSON is assumed.

Use `-table` to specify the name of a table to which the records are loaded. The alternative way to specify the table is to add the table specification "Table: <name>" before its records in the file.

For example, a file containing the records of 2 tables "users" and "email":

```
Table: users
<records of users>
...
Table: emails
<record of emails>
...
```

load

```
load -file <path to file>
```

Load the named file and interpret its contents as a script of commands to be executed. If any command in the script fails execution will end.

For example, suppose the following commands are collected in the script file `test.sql`:

```
### Begin Script ###
load -file test.ddl
import -table users -file users.json
### End Script ###
```

Where the file `test.ddl` would contain content like this:

```
DROP TABLE IF EXISTS users;
CREATE TABLE users(id INTEGER, firstname STRING, lastname STRING,
age INTEGER, primary key (id));
```

And the file `users.json` would contain content like this:

```
{ "id":1,"firstname":"Dean","lastname":"Morrison","age":51}
{ "id":2,"firstname":"Idona","lastname":"Roman","age":36}
{ "id":3,"firstname":"Bruno","lastname":"Nunez","age":49}
```

Then, the script can be run by using the `load` command in the shell:

```
> java -jar KVHOME/lib/sql.jar -helper-hosts node01:5000 \
-store kvstore
sql-> load -file ./test.sql
Statement completed successfully.
Statement completed successfully.
Loaded 3 rows to users.
```

mode

```
mode [COLUMN | LINE | JSON [-pretty] | CSV]
```

Sets the output mode of query results. The default value is JSON.

For example, a table shown in COLUMN mode:

```
sql-> mode column;
sql-> SELECT * from users;
+-----+-----+-----+-----+
| id | firstname | lastname | age |
+-----+-----+-----+-----+
| 8 | Len | Aguirre | 42 |
| 10 | Montana | Maldonado | 40 |
| 24 | Chandler | Oneal | 25 |
| 30 | Pascale | Mcdonald | 35 |
| 34 | Xanthus | Jensen | 55 |
| 35 | Ursula | Dudley | 32 |
| 39 | Alan | Chang | 40 |
| 6 | Lionel | Church | 30 |
| 25 | Alyssa | Guerrero | 43 |
| 33 | Gannon | Bray | 24 |
| 48 | Ramona | Bass | 43 |
| 76 | Maxwell | Mcleod | 26 |
| 82 | Regina | Tillman | 58 |
| 96 | Iola | Herring | 31 |
| 100 | Keane | Sherman | 23 |
+-----+-----+-----+-----+
...
```

100 rows returned

Empty strings are displayed as an empty cell.

```
sql-> mode column;
sql-> SELECT * from tabl where id = 1;
+-----+-----+-----+-----+
| id | s1 | s2 | s3 |
+-----+-----+-----+-----+
| 1 | NULL | | NULL |
+-----+-----+-----+-----+
```

1 row returned

For nested tables, indentation is used to indicate the nesting under column mode:

```
sql-> SELECT * from nested;
```

id	name	details	
1	one	address	
		city	Waitakere
		country	French Guiana
		zipcode	7229
		attributes	
		color	blue
		price	expensive
		size	large
		phone	[(08)2435-0742, (09)8083-8862, (08)0742-2526]
3	three	address	
		city	Viddalba
		country	Bhutan
		zipcode	280071
		attributes	
		color	blue
		price	cheap
		size	small
		phone	[(08)5361-2051, (03)5502-9721, (09)7962-8693]

...

For example, a table shown in LINE mode, where the result is displayed vertically and one value is shown per line:

```
sql-> mode line;
sql-> SELECT * from users;
```

> Row 1

id	8
firstname	Len
lastname	Aguirre
age	42

> Row 2

id	10
firstname	Montana
lastname	Maldonado
age	40

> Row 3

id	24
firstname	Chandler
lastname	Oneal
age	25

...

100 rows returned

Just as in COLUMN mode, empty strings are displayed as an empty cell:

```
sql-> mode line;
sql-> SELECT * from tabl where id = 1;
```

```
> Row 1
+-----+-----+
| id      | 1      |
| s1      | NULL   |
| s2      |        |
| s3      | NULL   |
+-----+-----+
```

1 row returned

For example, a table shown in JSON mode:

```
sql-> mode json;
sql-> SELECT * from users;
{"id":8,"firstname":"Len","lastname":"Aguirre","age":42}
{"id":10,"firstname":"Montana","lastname":"Maldonado","age":40}
{"id":24,"firstname":"Chandler","lastname":"Oneal","age":25}
{"id":30,"firstname":"Pascale","lastname":"Mcdonald","age":35}
{"id":34,"firstname":"Xanthus","lastname":"Jensen","age":55}
{"id":35,"firstname":"Ursula","lastname":"Dudley","age":32}
{"id":39,"firstname":"Alan","lastname":"Chang","age":40}
{"id":6,"firstname":"Lionel","lastname":"Church","age":30}
{"id":25,"firstname":"Alyssa","lastname":"Guerrero","age":43}
{"id":33,"firstname":"Gannon","lastname":"Bray","age":24}
{"id":48,"firstname":"Ramona","lastname":"Bass","age":43}
{"id":76,"firstname":"Maxwell","lastname":"Mcleod","age":26}
{"id":82,"firstname":"Regina","lastname":"Tillman","age":58}
{"id":96,"firstname":"Iola","lastname":"Herring","age":31}
{"id":100,"firstname":"Keane","lastname":"Sherman","age":23}
{"id":3,"firstname":"Bruno","lastname":"Nunez","age":49}
{"id":14,"firstname":"Thomas","lastname":"Wallace","age":48}
{"id":41,"firstname":"Vivien","lastname":"Hahn","age":47}
...
100 rows returned
```

Empty strings are displayed as "".

```
sql-> mode json;
sql-> SELECT * from tabl where id = 1;
{"id":1,"s1":null,"s2":"","s3":"NULL"}
```

1 row returned

Finally, a table shown in CSV mode:

```
sql-> mode csv;
sql-> SELECT * from users;
8,Len,Aguirre,42
10,Montana,Maldonado,40
24,Chandler,Oneal,25
30,Pascale,Mcdonald,35
34,Xanthus,Jensen,55
35,Ursula,Dudley,32
39,Alan,Chang,40
6,Lionel,Church,30
25,Alyssa,Guerrero,43
33,Gannon,Bray,24
48,Ramona,Bass,43
76,Maxwell,Mcleod,26
```

```
82,Regina,Tillman,58
96,Iola,Herring,31
100,Keane,Sherman,23
3,Bruno,Nunez,49
14,Thomas,Wallace,48
41,Vivien,Hahn,47
...
100 rows returned
```

Like in JSON mode, empty strings are displayed as "".

```
sql-> mode csv;
sql-> SELECT * from tabl where id = 1;
1,NULL,"","NULL"

1 row returned
```

Note:

Only rows that contain simple type values can be displayed in CSV format. Nested values are not supported.

output

```
output [stdout | file]
```

Enables or disables output of query results to a file. If no argument is specified, it shows the current output.

page

```
page [on | <n> | off]
```

Turns query output paging on or off. If specified, *n* is used as the page height.

If *n* is 0, or "on" is specified, the default page height is used. Setting *n* to "off" turns paging off.

show faults

```
show faults [-last] [-command <index>]
```

Encapsulates commands that display the state of the store and its components.

show query

```
show query <statement>
```

Displays the query plan for a query.

For example:

```
sql-> show query SELECT * from Users;
RECV([6], 0, 1, 2, 3, 4)
```

```
[
  DistributionKind : ALL_PARTITIONS,
  Number of Registers :7,
  Number of Iterators :12,
  SFW([6], 0, 1, 2, 3, 4)
  [
    FROM:
    BASE_TABLE([5], 0, 1, 2, 3, 4)
    [Users via primary index] as $$Users

    SELECT:
    *
  ]
]
```

show tables

```
show [as json] tables | table <table_name>
```

Shows either all tables currently existing in the store, or the named table.

show users

```
show [as json] users | user <user_name>
```

Shows either all the users currently existing in the store, or the named user.

show roles

```
show [as json] roles | role <role_name>
```

Shows either all the roles currently defined for the store, or the named role.

timeout

```
timeout [<timeout_ms>]
```

Configures or displays the request timeout for this session. If not specified, it shows the current value of request timeout.

timer

```
timer [on | off]
```

Turns the measurement and display of execution time for commands on or off. If not specified, it shows the current state of `timer`. For example:

```
sql-> timer on
sql-> SELECT * from users where id <= 10 ;
+----+-----+-----+-----+
| id | firstname | lastname | age |
+----+-----+-----+-----+
| 8 | Len | Aguirre | 42 |
| 10 | Montana | Maldonado | 40 |
| 6 | Lionel | Church | 30 |
| 3 | Bruno | Nunez | 49 |
| 2 | Idona | Roman | 36 |
```

	4		Cooper		Morgan		39	
	7		Hanae		Chapman		50	
	9		Julie		Taylor		38	
	1		Dean		Morrison		51	
	5		Troy		Stuart		30	
+-----+-----+-----+								

10 rows returned

Time: 0sec 98ms

verbose

verbose [on | off]

Toggles or sets the global verbosity setting. This property can also be set on a per-command basis using the `-verbose` flag.

version

version

Display client version information.