

Oracle® REST Data Services

SODA for REST Developer's Guide

Release 3.0

E58123-13

November 2016

Oracle REST Data Services SODA for REST Developer's Guide, Release 3.0

E58123-13

Copyright © 2014, 2016, Oracle and/or its affiliates. All rights reserved.

Primary Author: Drew Adams

Contributing Authors: Douglas McMahon, Sheila Moore, Maxim Orgiyan, Josh Spiegel

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	ix
Audience	ix
Documentation Accessibility	ix
Related Documents.....	ix
Conventions.....	ix
1 SODA for REST	
1.1 SODA for REST Overview.....	1-1
1.1.1 REST	1-2
1.1.2 Document Collections and Document-Centric APIs.....	1-2
1.1.3 SODA for REST Operations	1-3
1.2 Installing SODA for REST	1-3
1.3 Getting Started with SODA for REST	1-5
1.3.1 Creating a New Collection.....	1-6
1.3.2 Getting the List of Available Collections	1-7
1.3.3 Deleting a Collection.....	1-7
1.3.4 Inserting a Document into a Collection	1-8
1.3.5 Retrieving a Document from a Collection	1-9
1.3.6 Deleting a Document from a Collection.....	1-9
1.3.7 Bulk-Inserting Documents from a JSON Array	1-10
1.3.8 Listing the Documents in a Collection	1-11
1.3.9 Updating a Document in a Collection.....	1-13
1.3.10 Using a Filter Specification to Select Documents From a Collection.....	1-14
1.4 SODA for REST HTTP Operations.....	1-16
1.4.1 SODA for REST HTTP Operation URIs	1-16
1.4.2 SODA for REST HTTP Operation Response Bodies	1-17
1.4.3 GET schema.....	1-18
1.4.4 GET collection	1-19
1.4.5 GET object.....	1-22
1.4.6 PUT collection	1-23
1.4.7 PUT object.....	1-24
1.4.8 DELETE collection.....	1-24

1.4.9	DELETE object	1-25
1.4.10	POST object	1-26
1.4.11	POST query	1-27
1.4.12	POST array insert	1-28
1.4.13	POST bulk delete	1-29
1.5	Collection Specifications	1-30
1.5.1	Key Assignment Method.....	1-34
1.5.2	Versioning Method	1-35
1.6	Security.....	1-37
1.6.1	Authentication Mechanisms	1-37
1.6.2	Security Considerations for Development and Testing.....	1-38

Index

List of Examples

1-1	Response Body.....	1-18
1-2	Collection Specification.....	1-33

List of Tables

1-1	SODA for REST HTTP Operations.....	1-16
1-2	Fields That Can Appear in Response Bodies.....	1-17
1-3	Additional Response Body Fields for Operations that Return Objects.....	1-18
1-4	Relationship of GET collection Parameters to Mode and Links Array.....	1-22
1-5	Collection Specification Fields.....	1-30
1-6	Key Assignment Methods.....	1-34
1-7	Versioning Methods.....	1-35

Preface

This document explains how to use the Oracle SODA for REST API.

Audience

This document is intended for SODA for REST users.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

None

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

SODA for REST

The Oracle SODA for REST API is described, including how to install and use it.

Familiarity with the following can help you take best advantage of the information presented here:

- Oracle relational database management system (RDBMS)
- JavaScript Object Notation (JSON)
- Hypertext Transfer Protocol (HTTP)

See Also:

- *Oracle Database SODA for Java Developer's Guide*, which explains how to use the Java client API on which SODA for REST is built.
 - *Oracle Database JSON Developer's Guide* for information about JSON in Oracle Database
-

Topics

- [SODA for REST Overview](#) (page 1-1)
- [Installing SODA for REST](#) (page 1-3)
- [Getting Started with SODA for REST](#) (page 1-5)
- [SODA for REST HTTP Operations](#) (page 1-16)
- [Collection Specifications](#) (page 1-30)
- [Security](#) (page 1-37)

1.1 SODA for REST Overview

Simple Oracle Document Access (SODA, for short) lets you create and store collections of documents in Oracle Database, retrieve them, and query them, without needing to know Structured Query Language (SQL) or how the data in the documents is stored in the database.

SODA for REST is a REST API that provides SODA using the representational state transfer (REST) architectural style. You can use it to perform create, read, update, and delete (CRUD) operations on documents of any kind, and you can use it to query JSON documents.

Note:

Documents used with SODA for REST are limited to approximately 2 gigabytes.

Note:

Oracle recommends that you use AL32UTF8 (Unicode) for your database character set. Otherwise:

- Data can be altered during input, because of lossy conversion to the database character set.
 - Query-by-example (QBE) can return unpredictable results.
-
-

1.1.1 REST

The REST architectural style was used to define HTTP 1.1 and Uniform Resource Identifiers (URIs). A REST-based API strongly resembles the basic functionality provided by an HTTP server, and most REST-based systems are implemented using an HTTP client and an HTTP server.

A typical REST implementation maps create, read, update, and delete (CRUD) operations to HTTP verbs `POST`, `GET`, `PUT`, and `DELETE`, respectively.

A key characteristic of a REST-based system is that it is *stateless*: the server does not track or manage client object state. Each operation performed against a REST-based server is atomic; it is considered a transaction in its own right. In a typical REST-based system, many facilities that are taken for granted in an RDBMS environment, such as locking and concurrency control, are left to the application to manage.

A main advantage of a REST-based system is that its services can be used from almost any modern programming platform, including traditional programming languages (such as C, C#, C++, JAVA, and PL/SQL) and modern scripting languages (such as JavaScript, Perl, Python, and Ruby).

See Also:

Principled Design of the Modern Web Architecture, by Roy T. Fielding and Richard N. Taylor, at:

<http://www.ics.uci.edu/~taylor/documents/2002-REST-TOIT.pdf>

1.1.2 Document Collections and Document-Centric APIs

A REST-based application typically uses a document collection to persist state.

In a document collection, each document has a unique identifier. The identifier is typically assigned by the server when the document is created, but client-assigned identifiers can also be used. Document identifiers are **metadata**—data about the individual documents. Other metadata that a document collection can track for each document include the date and time that it was created and the date and time that it was last modified.

1.1.3 SODA for REST Operations

SODA for REST provides an HTTP-based service that implements the REST architectural style and allows documents to be stored in, and retrieved from, document collections managed by Oracle Database.

The API defines a set of simple operations that you can perform, including the following:

- List the set of available collections
- Create a collection
- Drop a collection
- Insert a document into a collection
- Retrieving a document from a collection
- Update a document in a collection
- Delete a document from a collection
- List the contents of a collection
- Search a collection

Your application can use the API operations to create and manipulate the JSON objects that it uses to persist application objects and state. To generate the JSON documents, your application can use JSON serialization techniques. When your application retrieves a document object, a JSON parser converts it to an application object.

See Also:

[SODA for REST HTTP Operations](#) (page 1-16) for detailed information about the operations defined by SODA for REST

1.2 Installing SODA for REST

Complete instructions are provided for installing SODA for REST.

To install SODA for REST:

1. Ensure that Oracle Database 12c Release 1 (12.1.0.2) with Merge Label Request (MLR) bundle patch 20885778 is installed. (Patch 20885778 obsoletes patch 20080249.)

Obtain the patch from My Oracle Support (<https://support.oracle.com>). Select tab **Patches & Updates**. Search for the patch number, 20885778 or access it directly at this URL: <https://support.oracle.com/rs?type=patch&id=20885778>.

2. Start the database.
3. Download Oracle REST Data Services (ORDS), and extract the zip file.

For instructions, see *Oracle REST Data Services Installation, Configuration, and Development Guide*.

4. Configure ORDS.

- If the database uses standard port 1521:

```
java -jar ords.war install
```

- If the database uses a nonstandard port (any port except 1521):

```
java -jar ords.war install advanced
```

Note:

When prompted:

- Do not skip the step of verifying/installing the Oracle REST Data Services schema.
 - Skip the steps that configure the PL/SQL Gateway.
 - Skip the steps that configure Application Express RESTful Services database users.
 - Decline to start the standalone server.
-
-

For more information, see *Oracle REST Data Services Installation, Configuration, and Development Guide*.

5. Connect to the schema that you want ORDS to access.**6. Enable ORDS in the schema by executing this SQL command:**

```
EXEC ords.enable_schema;  
COMMIT;
```

7. Grant role SODA_APP to the database schema (user account) *schema* that you enabled in step 6 (page 1-4):

```
GRANT SODA_APP TO schema;
```

8. Only if you are in a development environment:**a. Remove the default security constraints:**

```
BEGIN  
  ords.delete_privilege_mapping(  
    'oracle.soda.privilege.developer',  
    '/soda/*');  
  COMMIT;  
END;
```

This enables *anonymous* access to the service and is *not* recommended for production systems. For more information about security, see [Security](#) (page 1-37).

b. Start ORDS in standalone mode:

```
java -jar ords.war standalone
```

For more information, see *Oracle REST Data Services Installation, Configuration, and Development Guide*.

Note:

Disabling security and running ORDS in standalone mode is not recommended in production environments.

9. In a web browser, open:

```
http://localhost:8080/ords/schema/soda/latest/
```

Where *schema* is the lowercase name of the schema in which you enabled ORDS in step 6 (page 1-4). If the installation succeeded, you see:

```
{"items": [], "more": false}
```

See Also:

- [Security](#) (page 1-37)
 - *Oracle REST Data Services Installation, Configuration, and Development Guide*
-

1.3 Getting Started with SODA for REST

A step-by-step walkthrough is provided for the basic SODA for REST operations, using examples that you can run. The examples use command-line tool cURL to send REST requests to the server.

The examples assume that you started ORDS as instructed in [Installing SODA for REST](#) (page 1-3), enabling ORDS in *schema*.

Some examples also use the sample JSON documents included in the zip file that you downloaded in installation step 3 (page 1-3). They are in the directory `ORDS_HOME/examples/soda/getting-started`.

See Also:

- <http://curl.haxx.se/> for information about command-line tool cURL
 - *Oracle REST Data Services Installation, Configuration, and Development Guide*
-

Steps

1. [Creating a New Collection](#) (page 1-6)
2. [Getting the List of Available Collections](#) (page 1-7)
3. [Deleting a Collection](#) (page 1-7)
4. [Inserting a Document into a Collection](#) (page 1-8)
5. [Retrieving a Document from a Collection](#) (page 1-9)
6. [Deleting a Document from a Collection](#) (page 1-9)
7. [Bulk-Inserting Documents from a JSON Array](#) (page 1-10)

8. [Listing the Documents in a Collection](#) (page 1-11)
9. [Updating a Document in a Collection](#) (page 1-13)
10. [Using a Filter Specification to Select Documents From a Collection](#) (page 1-14)

1.3.1 Creating a New Collection

An example is given of creating a new collection.

To create a new collection, run this command, where *MyCollection* is the name of the collection:

```
curl -i -X PUT http://localhost:8080/ords/schema/soda/latest/MyCollection
```

The preceding command sends a PUT request with URL `http://localhost:8080/ords/schema/soda/latest/MyCollection`, to create a collection named *MyCollection*. The `-i` command-line option causes cURL to include the HTTP response headers in the output. If the operation succeeds then the output looks similar to this:

```
HTTP/1.1 201 Created
Cache-Control: private,must-revalidate,max-age=0
Location: http://localhost:8080/ords/schema/soda/latest/MyCollection/
Content-Length: 0
```

Response code 201 indicates that the operation succeeded. A PUT operation that results in the creation of a new collection—a PUT `collection` operation—returns no response body.

A successful PUT `collection` operation creates a database table to store the new collection. One way to see the details of this table is using SQL*Plus command `describe`:

```
SQL> describe "MyCollection"
Name                                         Null?    Type
-----
ID                                           NOT NULL VARCHAR2(255)
CREATED_ON                                  NOT NULL TIMESTAMP(6)
LAST_MODIFIED                               NOT NULL TIMESTAMP(6)
VERSION                                      NOT NULL VARCHAR2(255)
JSON_DOCUMENT                               BLOB
```

The preceding table reflects the default collection configuration. To create a custom collection configuration, provide a collection specification as the body of the PUT operation. For information about collection specifications, see [Collection Specifications](#) (page 1-30).

Caution:

To drop a collection, proceed as described in [Deleting a Collection](#) (page 1-7). Do *not* use SQL to drop the database table that underlies a collection. Collections have persisted metadata, in addition to the documents that are stored in the collection table.

See Also:

- [PUT collection](#) (page 1-23) for more information about this operation
 - [Getting the List of Available Collections](#) (page 1-7)
-
-

1.3.2 Getting the List of Available Collections

An example is given of listing available collections.

To obtain a list of the collections available in *schema*, run this command:

```
curl -X GET http://localhost:8080/ords/schema/soda/latest
```

That sends a GET request with the URL `http://localhost:8080/ords/schema/soda/latest` and returns this response body:

```
{
  "items" : [
    {
      "name": "MyCollection",
      "properties": {
        "schemaName": "SCHEMA",
        "tableName": "MyCollection",
        ...
      }
    }
  ],
  "links" : [
    {
      "rel" : "canonical",
      "href" : "http://localhost:8080/ords/schema/soda/latest/MyCollection"
    }
  ]
},
"more" : false
}
```

The response body includes all available collections in *schema*, which in this case is only collection *MyCollection*.

A successful GET *collection* operation returns response code 200, and the response body is a JSON object that contains an array of available collections and includes the collection specification for each collection.

See Also:

- [GET schema](#) (page 1-18)
 - [Collection Specifications](#) (page 1-30)
 - [Deleting a Collection](#) (page 1-7)
-
-

1.3.3 Deleting a Collection

An example is given of deleting a collection.

To delete *MyCollection*, run this command:

```
curl -i -X DELETE http://localhost:8080/ords/schema/soda/latest/MyCollection
```

The preceding command sends a DELETE request with the URL `http://localhost:8080/ords/schema/soda/latest/MyCollection` and returns:

```
HTTP/1.1 200 OK
Cache-Control: private,must-revalidate,max-age=0
Content-Length: 0
```

Response code 200 indicates that the operation succeeded. A DELETE operation that results in the deletion of a collection—a DELETE collection operation—returns no response body.

To verify that the collection was deleted, get the list of available collections in *schema*:

```
curl -X GET http://localhost:8080/ords/schema/soda/latest
```

If *MyCollection* was deleted, the preceding command returns:

```
{
  "items" : [],
  "more" : false
}
```

Create *MyCollection* again, so that you can use it in the next step:

```
curl -X PUT http://localhost:8080/ords/schema/soda/latest/MyCollection
```

See Also:

- [DELETE collection](#) (page 1-24)
 - [Inserting a Document into a Collection](#) (page 1-8)
-
-

1.3.4 Inserting a Document into a Collection

An example is given of inserting a document into a collection.

The example uses file `po.json`, which was included in the download. The file contains a JSON document that contains a purchase order. To load the JSON document into *MyCollection*, run this command:

```
curl -X POST --data-binary @po.json -H "Content-Type: application/json"
http://localhost:8080/ords/schema/soda/latest/MyCollection
```

The preceding command sends a POST request with the URL `http://localhost:8080/ords/schema/soda/latest/MyCollection`. It outputs something like this:

```
{
  "items" : [
    {
      "id" : "2FFD968C531C49B9A7EAC4398DFC02EE",
      "etag" : "C1354F27A5180FF7B828F01CBBC84022DCF5F7209DBF0E6DFFCC626E3B0400C3",
      "lastModified" : "2014-09-22T21:25:19.564394Z",
      "created" : "2014-09-22T21:25:19.564394Z"
    }
  ],
  "hasMore" : false,
  "count" : 1
}
```

A successful `POST object` operation returns response code 200. The response body is a JSON document that contains the identifier that the server assigned to the document when you inserted it into the collection, as well as the current ETag and last-modified time stamp for the inserted document.

Note:

If you intend to retrieve the document then copy the document identifier (the value of field `"id"`), to use for retrieval.

See Also:

- [POST object](#) (page 1-26) for more information about this operation
 - [Retrieving a Document from a Collection](#) (page 1-9)
-

1.3.5 Retrieving a Document from a Collection

An example is given of retrieving a document from a collection.

To retrieve the document that was inserted in [Inserting a Document into a Collection](#) (page 1-8), run this command, where `id` is the document identifier that you copied when inserting the document:

```
curl -X GET http://localhost:8080/ords/schema/soda/latest/MyCollection/id
```

A successful `GET document` operation returns response code 200. The response body contains the retrieved document.

If `id` does not exist in `MyCollection` then the response code is 404, as you can see by changing `id` to such an identifier:

```
curl -X GET http://localhost:8080/ords/schema/soda/latest/MyCollection/  
2FFD968C531C49B9A7EAC4398DFC02EF
```

```
{  
  "type" : "http://www.w3.org/Protocols/rfc2616/rfc2616-sec10.html#sec10.4.1",  
  "status" : 404,  
  "title" : "Key 2FFD968C531C49B9A7EAC4398DFC02EF not found in collection MyCollection.",  
  "o:errorCode" : "REST-02001"  
}
```

See Also:

- [GET object](#) (page 1-22)
 - [Deleting a Document from a Collection](#) (page 1-9)
-

1.3.6 Deleting a Document from a Collection

An example is given of deleting a document from a collection.

To delete, from `MyCollection`, the document that you retrieved in [Retrieving a Document from a Collection](#) (page 1-9), run this command (where `id` is the document identifier):

```
curl -i -X DELETE http://localhost:8080/ords/schema/soda/latest/MyCollection/id
```

The preceding command sends a DELETE request with URL `http://localhost:8080/ords/schema/soda/latest/MyCollection/id`, and it returns this:

```
HTTP/1.1 200 OK
Cache-Control: private,must-revalidate,max-age=0
Content-Length: 0
```

Response code 200 indicates that the operation succeeded. A DELETE operation that results in the deletion of an object from a collection—a DELETE object operation—returns no response body.

See Also:

- [DELETE object](#) (page 1-25)
 - [Bulk-Inserting Documents from a JSON Array](#) (page 1-10)
-
-

1.3.7 Bulk-Inserting Documents from a JSON Array

An example is given of bulk-inserting a set of documents into a collection from a JSON array of documents. The bulk insert operation is also called POST array insert.

This example uses file `POList.json`, which was included in the download. The file contains a JSON array of 70 purchase orders. To load the purchase orders into collection `MyCollection`, run this command:

```
curl -X POST --data-binary @POList.json -H "Content-Type: application/json"
http://localhost:8080/ords/schema/soda/latest/MyCollection?action=insert
```

Parameter value `action=insert` causes the array to be inserted as a set of documents, rather than as a single document.

The preceding command sends a POST request with the URL `http://localhost:8080/ords/schema/soda/latest/MyCollection`, and it outputs something like this:

```
{
  "items" : [
    {
      "id" : "6DEAF8F011FD43249E5F60A93B850AB9",
      "etag" : "49205D7E916EAED914465FCFF029B2795885A1914966E0AE82D4CCDBBE2EAF8E",
      "lastModified" : "2014-09-22T22:39:15.546435Z",
      "created" : "2014-09-22T22:39:15.546435Z"
    },
    {
      "id" : "C9FF7685D48E4E4B8641D8401ED0FB68",
      "etag" : "F3EB514BEDE6A6CC337ADA0F5BE6DEF5D451E68CE645729224BB6707FBE1F4F",
      "lastModified" : "2014-09-22T22:39:15.546435Z",
      "created" : "2014-09-22T22:39:15.546435Z"
    },
    ...
  ],
  "hasMore":false,
  "count":70
}
```

A successful `POST` array insert operation returns response code 200. The response body is a JSON document that contains the identifier, ETag, and last-modified time stamp for each inserted document.

Copy an "id" field value returned by your own `POST` array insert operation (not a value from the preceding example). Query the collection using `SQL*Plus` or `SQL Developer`, substituting your copied value for *identifier*:

```
SELECT json_value(json_document FORMAT JSON, '$.Reference')
       FROM "MyCollection" WHERE id = 'identifier';
```

```
JSON_VALUE(JSON_DOCUMENTFORMATJSON, '$.REFERENCE')
```

```
-----
MSULLIVA-20141102
```

Note:

In the SQL `SELECT` statement, you must specify the table name `MyCollection` as a quoted identifier, because it is mixed-case (the table name is the same as the collection name).

Because `MyCollection` has the default configuration, which stores the JSON document in a BLOB column, you must include `FORMAT JSON` when using the SQL/JSON function `json_value`. You cannot use the simplified JSON syntax.

See Also:

- [POST array insert](#) (page 1-28)
 - [Listing the Documents in a Collection](#) (page 1-11)
-
-

1.3.8 Listing the Documents in a Collection

An example is given of listing the documents in a collection, using a `GET` operation.

You can use parameters to control the result. For example, you can:

- Limit the number of documents returned
- Return only document identifiers (keys), only document contents, or both keys and contents
- Return a range of documents, based on keys or last-modified time stamps
- Specify the order of the list of returned documents

To list the documents in `MyCollection`, returning their keys and other metadata but not their content, run the following command.

```
curl -X GET http://localhost:8080/ords/schema/soda/latest/MyCollection?fields=id
```

The preceding command outputs something like this:

```
{
  "items" : [
    {
      "id" : "023C4A6581D84B71A5C0D5D364CE8484",
```

```
"etag": "3484DFB604DDA3FBC0C681C37972E7DD8C5F4457ACE32BD16960D4388C5A7C0E",
"lastModified" : "2014-09-22T22:39:15.546435Z",
"created": "2014-09-22T22:39:15.546435Z"
},
{
  "id" : "06DD0319148E40A7B8AA48E39E739184",
  "etag" : "A19A1E9A3A38B1BAE3EE52B93350FBD76309CBFC4072A2BECD95BCA44D4849DD",
  "lastModified" : "2014-09-22T22:39:15.546435Z",
  "created" : "2014-09-22T22:39:15.546435Z"
},
...
],
"hasMore" : false,
"count" : 70,
"offset":0,
"limit":100,
"totalResults":70
}
```

A successful GET collection operation returns response code 200, and the response body is a JSON document that lists the documents in the collection. If the collection is empty, the response body is an empty items array.

To list at most 10 documents in MyCollection, returning their keys, content, and other metadata, run this command:

```
curl -X GET "http://localhost:8080/ords/schema/soda/latest/MyCollection?fields=all&limit=10"
```

The preceding command outputs something like this:

```
{
  "items": [ ... ],
  "hasMore" : true,
  "count" : 10,
  "offset" : 0,
  "limit" : 10,
  "links" : [{
    "rel" : "next",
    "href" :
"http://localhost:8080/ords/schema/soda/latest/MyCollection?offset=10&limit=10"
  }]
}
```

Note:

Including document content makes the response body much larger. Oracle recommends including the content in the response body only if you will need the content later. Retrieving the content from the response body is more efficient than retrieving it from the server.

The metadata in the response body shows that 10 documents were requested ("limit" : 10) and 10 documents were returned ("count" : 10), and that more documents are available ("hasMore" : true). To fetch the next set of documents, you can use the URL in the field "links". "href".

The maximum number of documents returned from a collection by the server is controlled by the following:

- URL parameter limit

- Configuration parameters `soda.maxLimit` and `soda.defaultLimit`

Note:

If you intend to update the document then copy the document identifier (value of field "id"), to use for updating.

See Also:

- [GET collection](#) (page 1-19)
 - [Updating a Document in a Collection](#) (page 1-13)
 - [Using a Filter Specification to Select Documents From a Collection](#) (page 1-14), which lets you list documents based on content
 - *Oracle REST Data Services Installation, Configuration, and Development Guide* for information about configuration parameters `soda.maxLimit` and `soda.defaultLimit`
-
-

1.3.9 Updating a Document in a Collection

An example is given of updating a document in a collection, that is, replacing it with a newer version. For this, you use a PUT operation.

The behavior of the PUT operation for a nonexistent document depends on the key-assignment method used by the collection.

- If the collection uses *server-assigned keys* (as does collection `MyCollection`) then an error is raised if you try to update a nonexistent document (that is, you specify a key that does not belong to any document in the collection).
- If the collection uses *client-assigned keys*, then trying to update a nonexistent document *inserts* into the collection a new document with the specified key.

Retrieve a document from `MyCollection` by running this command, where `id` is the document identifier that you copied in [Listing the Documents in a Collection](#) (page 1-11):

```
curl -X GET http://localhost:8080/ords/schema/soda/latest/MyCollection/id
```

The preceding command outputs the retrieved document.

To update this document with the content of file `poUpdated.json`, which was included in the download, execute this command:

```
curl -i -X PUT --data-binary @poUpdated.json -H "Content-Type: application/json"
http://localhost:8080/ords/schema/soda/latest/MyCollection/id
```

The preceding command outputs something like this:

```
HTTP/1.1 200 OK
Cache-Control: no-cache,must-revalidate,no-store,max-age=0
ETag: A0B07E0A6D000358C546DC5D8D5059D9CB548A1A5F6F2CAD66E2180B579CCB6D
Last-Modified: Mon, 22 Sep 2014 16:42:35 PDT
Location: http://localhost:8080/ords/schema/soda/latest/MyCollection/
023C4A6581D84B71A5C0D5D364CE8484/
Content-Length: 0
```

The response code 200 indicates that the operation succeeded. A `PUT` operation that results in the successful update of a document in a collection—a `PUT object` operation—returns no response body.

To verify the document has been updated, rerun this command:

```
curl -X GET http://localhost:8080/ords/schema/soda/latest/MyCollection/id
```

The preceding command returns:

```
{
  "PONumber": 1,
  "Content" : "This document has been updated...."
}
```

See Also:

- [PUT object](#) (page 1-24)
 - [Key Assignment Method](#) (page 1-34)
 - [Using a Filter Specification to Select Documents From a Collection](#) (page 1-14)
-

1.3.10 Using a Filter Specification to Select Documents From a Collection

Examples are given of using a filter specification, or query-by-example (QBE), to define query criteria for selecting documents from a collection.

The examples use the `QBE.*.json` files that are included in the zip file that you downloaded in installation step 3 (page 1-3). They are in directory `ORDS_HOME/examples/soda/getting-started`.

See Also:

- [POST query](#) (page 1-27)
 - *Oracle Database SODA for Java Developer's Guide* for information about filter specifications and QBE
-

1.3.10.1 QBE.1.json

The query-by-example (QBE) in file `QBE.1.json` returns a list of nine documents, each of which has "TGATES" as the value of field `User`.

This is the query in file `QBE.1.json`:

```
{ "User" : "TGATES" }
```

To execute the query, run this command:

```
curl -X POST --data-binary @QBE.1.json -H "Content-Type: application/json"
http://localhost:8080/ords/schema/soda/latest/MyCollection?action=query
```

A successful `POST query` operation returns response code 200 and a list of documents that satisfy the query criteria.

Because the command has no `fields` parameter, the default value `fields=all` applies, and the response body contains both the metadata and the content of each document.

Note:

Including document content makes the response body much larger. Oracle recommends including the content in the response body only if you need the content for a subsequent operation. Retrieving the content from the response body is more efficient than retrieving it from the server.

To execute the queries in the other `QBE.*.json` files, run commands similar to the preceding one.

1.3.10.2 QBE.2.json

The query-by-example (QBE) in file `QBE.2.json` selects documents where the value of field `UPCCode` equals "13023015692". `UPCCode` is a field of object `Part`, which is a field of array `LineItems`. Because no array offset is specified for `LineItems`, the query searches all elements of the array.

This is the query in file `QBE.2.json`:

```
{ "LineItems.Part.UPCCode" : "13023015692" }
```

Note:

Keyword `"$eq"` in the query is implied. See *Oracle Database SODA for Java Developer's Guide* for more information.

1.3.10.3 QBE.3.json

The query-by-example (QBE) in file `QBE.3.json` selects documents where the value of field `ItemNumber`, in object `LineItems`, is greater than 4. Keyword `"$gt"` is required.

This is the query in file `QBE.3.json`:

```
{ "LineItems.ItemNumber" : { "$gt" : 4 } }
```

1.3.10.4 QBE.4.json

The query-by-example (QBE) in file `QBE.4.json` selects documents where the value of field `UPCCode` equals "13023015692" and the value of field `ItemNumber` equals 3. Keyword `"$and"` is optional.

This is the query in file `QBE.4.json`:

```
{ "$and" : [
  { "LineItems.Part.UPCCode" : "13023015692" },
  { "LineItems.ItemNumber" : 3 }
]
}
```

See Also:

Oracle Database SODA for Java Developer's Guide

1.4 SODA for REST HTTP Operations

The SODA for REST HTTP operations are described.

[Table 1-1](#) (page 1-16) summarizes the HTTP operations that SODA for REST provides. For complete descriptions of the operations, click the links in the left column.

Table 1-1 SODA for REST HTTP Operations

Operation	Description
GET schema (page 1-18)	Gets some or all collection names in a schema.
GET collection (page 1-19)	Gets all or a subset of objects from a collection, using parameters to specify the subset. You can page through the return set.
GET object (page 1-22)	Gets a specified object from a collection.
PUT collection (page 1-23)	Creates a collection if it does not exist.
PUT object (page 1-24)	Replaces a specified object with an uploaded object (typically a new version). If the collection has client-assigned keys and the uploaded object is not already in the collection, then PUT inserts the uploaded object into the collection.
DELETE collection (page 1-24)	Deletes a collection.
DELETE object (page 1-25)	Deletes a specified object from a collection.
POST object (page 1-26)	Puts an uploaded object in a specified collection, assigning and returning its key. Collection must use server-assigned keys.
POST query (page 1-27)	Gets all or a subset of objects from a collection, using a filter to specify the subset. You cannot page through the return set.
POST array insert (page 1-28)	Inserts an array of objects into a specified collection, assigning and returning their keys.
POST bulk delete (page 1-29)	Deletes all or a subset of objects from a collection.

1.4.1 SODA for REST HTTP Operation URIs

A SODA for REST HTTP operation is specified by a Universal Resource Identifier (URI).

The URI has this form:

```
/ords/schema/soda/[version/[collection/[{key/|?action=action}]]]
```

where:

- `ords` is the directory of the Oracle REST Data Services (ORDS) listener, of which SODA for REST is a component.
- `schema` is the name of an Oracle Database schema that has been configured as an end point for SODA for REST.
- `soda` is the name given to the Oracle Database JSON service when mapped as a template within ORDS.
- `version` is the version number of `soda`.
- `collection` is the name of a set of objects stored in `schema`.

Typically, an object is a JSON document, but it can be a Multipurpose Internet Mail Extensions (MIME) type (for example, image, audio, or video).

A JSON document is represented as textual JSON.

Typically, an application uses a collection to hold all instances of a particular type of object. Thus, a collection is roughly analogous to a table in a relational database. One column stores keys and another column stores content.

- `key` is a string that uniquely identifies an object in `collection`.
A **specified object** is specified by its key.
- `action` is either `query`, `index`, `unindex`, `insert`, `update`, or `delete`.

1.4.2 SODA for REST HTTP Operation Response Bodies

If a SODA for REST HTTP operation returns information or objects, it does so in a response body.

For the operation [GET object](#) (page 1-22), the response body is a single object.

[Table 1-2](#) (page 1-17) lists and describes fields that can appear in response bodies.

Table 1-2 Fields That Can Appear in Response Bodies

Field	Description
<code>key</code>	String that uniquely identifies an object (typically a JSON document) in a collection.
<code>etag</code>	HTTP entity tag (ETag)—checksum or version.
<code>created</code>	Created-on time stamp.
<code>lastModified</code>	Last-modified time stamp.
<code>value</code>	Object contents (applies only to JSON object).
<code>mediaType</code>	HTTP Content-Type (applies only to non-JSON object).
<code>bytes</code>	HTTP Content-Length (applies only to non-JSON object).
<code>items</code>	List of one or more collections or objects that the operation found or created. This field can be followed by the fields in Table 1-3 (page 1-18).

If an operation creates or returns objects, then its response body can have the additional fields in [Table 1-3](#) (page 1-18). The additional fields appear after field `items`.

Table 1-3 Additional Response Body Fields for Operations that Return Objects

Field	Description
name	Name of collection. This field appears only in the response body of GET schema (page 1-18).
properties	Properties of collection. This field appears only in the response body of GET schema (page 1-18).
hasMore	true if <code>limit</code> was reached before available objects were exhausted, false otherwise. This field is always present.
limit	Server-imposed maximum collection (row) limit.
offset	Offset of first object returned (if known).
count	Number of objects returned. This is the only field that can appear in the response body of POST bulk delete (page 1-29).
totalResults	Number of objects in collection (if requested)
links	Possible final field for <code>GET collection</code> operation. For details, see GET collection (page 1-19).

[Example 1-1](#) (page 1-18) shows the structure of a response body that returns 25 objects. The first object is a JSON object and the second is a jpeg image. The collection that contains these objects contains additional objects.

Example 1-1 Response Body

```
{
  "items" : [
    {
      "id"       : "key_of_object_1",
      "etag"     : "etag_of_object_1",
      "lastModified" : "lastmodified_timestamp_of_object_1",
      "value"    : {object_1}
    },
    {
      "id"       : "key_of_object_2",
      "etag"     : "etag_of_object_2",
      "lastModified" : "lastmodified_timestamp_of_object_2",
      "mediaType"  : "image/jpeg",
      "bytes"     : 1234
    },
    ...
  ],
  "hasMore" : true,
  "limit"   : 100,
  "offset"  : 50,
  "count"   : 25,
  "links"   : [ ... ]
}
```

1.4.3 GET schema

GET schema gets all or a subset of collection names in a schema.

See Also:[Listing the Documents in a Collection](#) (page 1-11)

1.4.3.1 URL Pattern for GET schema

The URL pattern for GET schema is described.

```
/ords/schema/soda/version/
```

Without parameters, GET *schema* gets all collection names in *schema*.

Parameter	Description
<code>limit=<i>n</i></code>	Limits number of listed collection names to <i>n</i> .
<code>fromID=<i>collection</i></code>	Starts getting with <i>collection</i> (inclusive).

1.4.3.2 Response Codes for GET schema

The response codes for GET schema are described.

200

Success—response body contains names and properties of collections in schema, ordered by name. For example:

```
{
  "items" : [
    { "name"      : "employees",
      "properties" : {...} },
    { "name"      : "departments",
      "properties" : {...} },
    ...
  ],
  "hasMore" : false
}
```

If `hasMore` is `true`, then to get the next batch of collection names specify `fromID=last_returned_collection`. (In the preceding example, *last_returned_collection* is "regions").

404

Either the schema was not found or access is not authorized.

1.4.4 GET collection

GET collection gets all or a subset of objects from a collection, using parameters to specify the subset. You can page through the set of returned objects.

See Also:

- [POST query](#) (page 1-27), which gets all or a subset of objects from a collection, using a filter instead of parameters. You *cannot* page through the set of returned objects.
- [Listing the Documents in a Collection](#) (page 1-11)

1.4.4.1 URL Pattern for GET collection

The URL pattern for GET collection is described.

/ords/schema/soda/version/collection/

Without parameters, GET `collection` gets all objects (both key and content) from `collection` and does not return the number of objects in `collection`.

Note:

For non-JSON objects in the collection, GET `collection` returns, instead of content, media type and (if known) size in bytes.

Parameter	Description
<code>limit=n</code>	Limits number of objects to <i>n</i> .
<code>offset=n</code>	Skips <i>n</i> objects before getting first object.
<code>fields={id value all}</code>	Gets only object <code>id</code> fields (keys), only object <code>value</code> fields (content), or <code>all</code> fields (both key and content). Regardless of the <code>fields</code> value, GET <code>collection</code> returns the other metadata that the collection stores for each document.
<code>totalResults=true</code>	Returns number of objects in collection. Note: Inefficient
<code>fromID=key</code>	Starts getting objects after <i>key</i> , in ascending order.
<code>toID=key</code>	Stops getting objects before <i>key</i> , in descending order.
<code>after=key</code>	Starts getting objects after <i>key</i> , in ascending order.
<code>before=key</code>	Stops getting objects before <i>key</i> , in descending order.
<code>since=timestamp</code>	Gets only objects with time stamp later than <i>timestamp</i> .
<code>until=timestamp</code>	Gets only objects with time stamp earlier than <i>timestamp</i> .

1.4.4.2 Response Codes for GET collection

The response codes for GET collection are described.

200

Success—response body contains the specified objects from `collection` (or only their keys, if you specified `fields=id`). For example:

```

{
  "items" : [
    {
      "id"          : "key_of_object_1",
      "etag"        : "etag_of_object_1",
      "lastModified" : "lastmodified_timestamp_of_object_1",
      "value"       : {object_1}
    },
    {
      "id"          : "key_of_object_2",
      "etag"        : "etag_of_object_2",
      "lastModified" : "lastmodified_timestamp_of_object_2",
      "value"       : {object_2}
    },
    {
      "id"          : "key_of_object_3",
      "etag"        : "etag_of_object_3",
      "lastModified" : "lastmodified_timestamp_of_object_3",
      "mediaType"   : "image/jpeg",
      "bytes"       : 1234
    },
    ...
  ],
  "hasMore" : true,
  "limit"   : 100,
  "offset"  : 50,
  "count"   : 25
  "links"   : [ ... ]
}

```

If `hasMore` is `true`, then to get the next batch of objects repeat the operation with an appropriate parameter. For example:

- `offset=n` if the response body includes the offset
- `toID=last_returned_key` or `before=last_returned_key` if the response body includes `descending=true`
- `fromID=last_returned_key` or `after=last_returned_key` if the response body does not include `descending=true`

For information about links, see [Links Array for GET collection](#) (page 1-21).

401

Read access to collection is not authorized.

404

Collection was not found.

1.4.4.3 Links Array for GET collection

The links array for GET collection is described.

The existence and content of the links array depends on the **mode** of the GET collection operation, which is determined by its parameters.

When the links array exists, it has an element for each returned object. Each element contains links from that object to other objects. The possible links are:

- `first`, which links the object to the first object in the collection
- `prev`, which links the object to the previous object in the collection
- `next`, which links the object to the next object in the collection

Using `prev` and `next` links, you can page through the set of returned objects.

[Table 1-4](#) (page 1-22) shows how `GET collection` parameters determine mode and the existence and content of the `links` array.

Table 1-4 Relationship of GET collection Parameters to Mode and Links Array

Parameter	Mode	Links Array
<code>fields=id</code>	Keys-only	Does not exist (regardless of other parameters).
<code>offset=n</code>	Offset	Has an element for each returned object. Each element has these links, except as noted: <ul style="list-style-type: none"> • <code>first</code> (except for first object) • <code>prev</code> (except for first object) • <code>next</code> (except for last object)
<code>fromID=key</code> <code>toID=key</code> <code>after=key</code> <code>before=key</code>	Keyed	Has an element for each returned object. Each element has these links, except as noted: <ul style="list-style-type: none"> • <code>prev</code> (except for first object) • <code>next</code> (except for last object)
<code>since=timestamp</code> <code>until=timestamp</code>	Timestamp	Does not exist.

1.4.5 GET object

`GET object` gets a specified object from a specified collection.

See Also:

[Retrieving a Document from a Collection](#) (page 1-9)

1.4.5.1 URL Pattern for GET object

The URL pattern for `GET object` is described.

`/ords/schema/soda/version/collection/key/`

No parameters.

1.4.5.2 Request Headers for GET object

The request headers for `GET object` are described.

Operation `GET object` accepts these optional request headers:

Header	Description
<code>If-Modified-Since=timestamp</code>	Returns response code 304 if object has not changed since <code>timestamp</code> .

Header	Description
If-None-Match	Returns response code 304 if object etag matches the current checksum or version.

1.4.5.3 Response Codes for GET object

The response codes for GET object are described.

200

Success—response body contains object identified by the URL pattern.

204

Object content is null.

304

Either object was not modified since modification date or checksum matches object etag (see [Request Headers for GET object](#) (page 1-22)).

401

Read access to collection or object is not authorized.

404

Collection or object was not found.

1.4.6 PUT collection

PUT collection creates a collection if it does not exist.

See Also:

[Creating a New Collection](#) (page 1-6)

1.4.6.1 URL Pattern for PUT collection

The URL pattern for PUT collection is described.

*/ords/schema/soda/version/**collection**/*

No parameters.

1.4.6.2 Request Body for PUT collection (Optional)

The request body for PUT collection is described.

See [Collection Specifications](#) (page 1-30).

1.4.6.3 Response Codes for PUT collection

The response codes for PUT collection are described.

200

Collection with the same name and properties already exists.

201

Success—collection was created.

401

Collection creation is not authorized.

1.4.7 PUT object

PUT object replaces a specified object in a specified collection with an uploaded object (typically a new version). If the collection has client-assigned keys and the uploaded object is not already in the collection, then PUT inserts the uploaded object into the collection.

See Also:

[Updating a Document in a Collection](#) (page 1-13)

1.4.7.1 URL Pattern for PUT object

The URL pattern for PUT object is described.

`/ords/schema/soda/version/collection/key/`

No parameters.

1.4.7.2 Request Body for PUT object

The request body for PUT object is the uploaded object.

1.4.7.3 Response Codes for PUT object

The response codes for PUT object are described.

200

Success—object was replaced.

401

Updating collection is not authorized.

405

Collection is read-only.

1.4.8 DELETE collection

DELETE collection deletes a collection.

To delete all *objects* from a collection, but *not delete the collection* itself, use [POST bulk delete](#) (page 1-29).

See Also:[Deleting a Collection](#) (page 1-7)

1.4.8.1 URL Pattern for DELETE collection

The URL pattern for DELETE collection is described.

/ords/schema/soda/version/collection/

No parameters.

1.4.8.2 Response Codes for DELETE collection

The response codes for DELETE collection are described.

200

Success—collection was deleted.

401

Deleting collection is not authorized.

404

Collection was not found.

1.4.9 DELETE object

DELETE object deletes a specified object from a specified collection.

See Also:[Deleting a Document from a Collection](#) (page 1-9)

1.4.9.1 URL Pattern for DELETE object

The URL pattern for DELETE object is described.

/ords/schema/soda/version/collection/key/

No parameters.

1.4.9.2 Response Codes for DELETE object

The response codes for DELETE object are described.

200

Success—object was deleted.

401

Either deleting from collection or deleting this object is not authorized.

404

Object was not found.

405

Collection is read-only.

1.4.10 POST object

POST object inserts an uploaded object into a specified collection, assigning and returning its key. The collection must use server-assigned keys.

If the collection uses client-assigned keys, use [PUT object](#) (page 1-24). For information about key assignment methods, see [Key Assignment Method](#) (page 1-34).

See Also:

[Inserting a Document into a Collection](#) (page 1-8)

1.4.10.1 URL Pattern for POST object

The URL pattern for POST object is described.

*/ords/schema/soda/version/**collection**/*

No parameters.

1.4.10.2 Request Body for POST object

The request body for POST object is the uploaded object to be inserted in the collection.

1.4.10.3 Response Codes for POST object

The response codes for POST object are described.

201

Success—object is in collection; response body contains server-assigned key and possibly other information. For example:

```
{
  "items" : [
    {
      "id"          : "key",
      "etag"        : "etag",
      "lastModified" : "timestamp"
      "created"     : "timestamp"
    }
  ],
  "hasMore" : false
}
```

202

Object was accepted and queued for asynchronous insertion; response body contains server-assigned key.

401

Inserting into collection is not authorized.

405

Collection is read-only.

501

Unsupported operation (for example, no server-side key assignment).

1.4.11 POST query

POST query gets all or a subset of objects from a collection, using a filter to specify the subset. You cannot page through the set of returned objects.

See Also:

- [GET collection](#) (page 1-19), which gets all objects, or a subset of these, from a collection, using parameters instead of a filter. You can page through the set of returned objects.
- [Using a Filter Specification to Select Documents From a Collection](#) (page 1-14)

1.4.11.1 URL Pattern for POST query

The URL pattern for POST query is described.

`/ords/schema/soda/version/collection?action=query`

Parameters are optional except as noted.

Parameter	Description
<code>action=query</code>	Required. Specifies kind of action.
<code>limit=n</code>	Limit number of returned objects to <i>n</i> .
<code>offset=n</code>	Skip <i>n</i> objects before returning objects.
<code>fields={id value all}</code>	Return object <code>id</code> (key) only, object <code>value</code> (content) only, or <code>all</code> (object key and content). Default: <code>all</code>

1.4.11.2 Request Body for POST query

If you omit the filter specification object from the request body of POST query then the operation gets all objects in the collection.

See Also:

Oracle Database SODA for Java Developer's Guide for information about SODA filter specifications.

1.4.11.3 Response Codes for POST query

The response codes for POST query are described.

200

Success—object is in collection; response body contains all objects in collection that match filter.

404

Either collection was not found or read access to collection is not authorized.

1.4.12 POST array insert

POST array insert inserts an array of objects into a specified collection, assigning and returning their keys.

See Also:

[Bulk-Inserting Documents from a JSON Array](#) (page 1-10)

1.4.12.1 URL Pattern for POST array insert

The URL pattern for POST array insert is described.

`/ords/schema/soda/version/collection?action=insert`

Parameter	Description
<code>action=insert</code>	Required. Specifies kind of action.

1.4.12.2 Request Body for POST array insert

The request body for POST array insert is an array of objects.

Array of objects.

1.4.12.3 Response Codes for POST array insert

The response codes for POST array insert are described.

200

Success—response body contains an array with the assigned keys for inserted objects. For example:

```
{
  "items" : [
    {
      "id"          : "12345678",
      "etag"       : "...",
      "lastModified" : "...",
      "created"    : "..."
    },
    {
      "id"          : "23456789",
      "etag"       : "...",

```

```

        "lastModified" : "...",
        "created"      : "...",
      }
    ],
    "hasMore" : false
  }

```

401

Inserting into collection is not authorized.

404

Collection was not found.

405

Collection is read-only.

1.4.13 POST bulk delete

POST bulk delete deletes all or a subset of objects from a specified collection, using a filter to specify the subset.

Note:

If you delete all objects from the collection, the empty collection continues to exist. To delete the collection itself, use [DELETE collection](#) (page 1-24).

1.4.13.1 URL Pattern for POST bulk delete

The URL pattern for POST bulk delete is described.

Either of the following:

```
/ords/schema/soda/version/collection?action=delete
```

```
/ords/schema/soda/version/collection?action=truncate
```

Parameter	Description
action=delete	Required. Specifies the deletion of all or a subset of objects from <i>collection</i> , using an optional filter to specify the subset. See the following warning.
action=truncate	Required. Specifies the deletion of all objects from <i>collection</i> . Does not use a filter.

Warning:

If you specify `action=delete` and omit the filter specification, or if the filter specification is empty, then the operation deletes all objects from the collection.

1.4.13.2 Request Body for POST bulk delete (Optional)

See *Oracle Database SODA for Java Developer's Guide* for information about SODA filter specifications.

1.4.13.3 Response Codes for POST bulk delete

The response codes for POST bulk delete are described.

200

Success—response body contains number of deleted collections. For example:

```
{
  "count" : 42
}
```

401

Deleting from collection is not authorized.

405

Collection is read-only.

1.5 Collection Specifications

A collection specification provides information about the Oracle Database table or view underlying the collection object. The table or view is created when you create the collection.

Note:

In collection specifications, you must use *strict* JSON syntax. That is, you must enclose each nonnumeric value in double quotation marks.

[Table 1-5](#) (page 1-30) describes the collection specification fields and their possible values.

Note:

If you omit one of the optional columns (created-on timestamp, last-modified timestamp, version, or media type) from the collection specification then no such column is created. At a minimum, a collection has a key column and a content column.

Table 1-5 Collection Specification Fields

Field	Description	Possible Values
schemaName	SQL name of schema that owns table or view underlying collection object.	—
tableName or viewName	SQL name of table or view underlying collection object.	—

Table 1-5 (Cont.) Collection Specification Fields

Field	Description	Possible Values
<code>keyColumn.name</code>	Name of key column.	Default: ID
<code>keyColumn.sqlType</code>	SQL data type of key column.	VARCHAR2 (default), NUMBER, RAW
<code>keyColumn.maxLength</code>	Maximum length of key column, if not of NUMBER data type.	Default: 255
<code>keyColumn.assignmentMethod</code>	Key assignment method.	SEQUENCE, GUID, UUID (default), or CLIENT
<code>keyColumn.sequenceName</code>	If <code>keyColumn.assignmentMethod</code> is SEQUENCE, then this field must specify the name of a database sequence.	Name of existing database sequence
<code>contentColumn.name</code>	Name of content column.	Default: JSON_DOCUMENT
<code>contentColumn.sqlType</code>	SQL data type of content column.	VARCHAR2, BLOB (default), CLOB
<code>contentColumn.maxLength</code>	Maximum length of content column, if not of LOB data type.	The default length is 4000 bytes. If <code>MAX_STRING_SIZE = STANDARD</code> then <code>maxLength</code> can be at most 4000 (bytes). If <code>MAX_STRING_SIZE = EXTENDED</code> , then <code>maxLength</code> can be at most 32767 (bytes).

Table 1-5 (Cont.) Collection Specification Fields

Field	Description	Possible Values
<code>contentColumn.validation</code>	<p>Validation level of content column. Corresponds to SQL condition <code>is json</code>, which determines the syntax to which JSON content must conform.</p> <p><code>STANDARD</code> validates according to the JSON RFC 4627 standard. (It corresponds to the strict syntax defined for Oracle SQL condition <code>is json</code>.)</p> <p><code>STRICT</code> is the same as <code>STANDARD</code>, except that it also verifies that the document does not contain duplicate JSON field names. (It corresponds to the strict syntax defined for Oracle SQL condition <code>is json</code> when the keywords <code>WITH UNIQUE KEYS</code> are also used.)</p> <p><code>LAX</code> validates more loosely. (It corresponds to the lax syntax defined for Oracle SQL condition <code>is json</code>.)</p> <p>Some of the relaxations that <code>LAX</code> allows include the following:</p> <ul style="list-style-type: none"> • It does not require JSON field names to be enclosed in double quotation marks ("). • It allows uppercase, lowercase, and mixed case versions of <code>true</code>, <code>false</code>, and <code>null</code>. • Numerals can be represented in additional ways. 	<code>STANDARD</code> (default), <code>STRICT</code> , <code>LAX</code>
<code>contentColumn.compress</code>	Compression level for SecureFiles stored in content column.	<code>NONE</code> (default), <code>HIGH</code> , <code>MEDIUM</code> , <code>LOW</code>
<code>contentColumn.cache</code>	Caching of SecureFiles stored in content column.	<code>TRUE</code> , <code>FALSE</code> (default)
<code>contentColumn.encrypt</code>	Encryption algorithm for SecureFiles stored in content column. ¹	<code>NONE</code> (default), <code>3DES168</code> , <code>AES128</code> , <code>AES192</code> , <code>AES256</code>
<code>creationTimeColumn.name</code>	<p>Name of optional created-on timestamp column.</p> <p>This column has SQL data type <code>TIMESTAMP</code> and default value <code>SYSTIMESTAMP</code>.</p>	Default: <code>CREATED_ON</code>
<code>lastModifiedColumn.name</code>	<p>Name of optional last-modified timestamp column.</p> <p>This column has SQL data type <code>TIMESTAMP</code> and default value <code>SYSTIMESTAMP</code>.</p>	Default: <code>LAST_MODIFIED</code>
<code>lastModifiedColumn.index</code>	Name of nonunique index on timestamp column. The index is created if a name is specified.	

Table 1-5 (Cont.) Collection Specification Fields

Field	Description	Possible Values
<code>versionColumn.name</code>	Name of optional version (ETag) column. This column has SQL data type <code>VARCHAR2 (255)</code> unless the method is <code>SEQUENTIAL</code> or <code>TIMESTAMP</code> , in which case it has data type <code>NUMBER</code> . Note: If the method is <code>TIMESTAMP</code> then the version is stored as an integer representation of the date and time with microsecond precision. It does not store a date/time string or a SQL date/time type.	Default: <code>VERSION</code>
<code>versionColumn.method</code>	Versioning method.	<code>SEQUENTIAL</code> , <code>TIMESTAMP</code> , <code>UUID</code> , <code>SHA256</code> (default), <code>MD5</code> , <code>NONE</code>
<code>mediaTypeColumn.name</code>	Name of optional object media type column. This column has SQL data type <code>VARCHAR2 (255)</code> .	
<code>readOnly</code>	Read/write policy: <code>TRUE</code> means read-only.	<code>TRUE</code> , <code>FALSE</code> (default)

¹ Set up Encryption Wallet before creating a collection with SecureFile encryption. For information about the `SET ENCRYPTION WALLET` clause of the `ALTER SYSTEM` statement, see *Oracle Database SQL Language Reference*.

[Example 1-2](#) (page 1-33) is a collection specification for an object whose underlying table is `HR.EMPLOYEES`.

Example 1-2 Collection Specification

```
{
  "schemaName"      : "HR",
  "tableName"       : "EMPLOYEES",
  "contentColumn"   :
  {
    "name"           : "EMP_DOC",
    "sqlType"        : "VARCHAR2",
    "maxLength"     : 4000,
    "validation"    : "STRICT",
    "compress"      : "HIGH",
    "cache"          : true,
    "encrypt"        : "AES128",
  },
  "keyColumn"       :
  {
    "name"           : "EMP_ID",
    "sqlType"        : "NUMBER",
    "assignmentMethod" : "SEQUENCE",
    "sequenceName"   : "EMPLOYEE_ID_SEQ"
  },
  "creationTimeColumn" :
  {
```

```

    "name"          : "CREATED_ON"
  },
  "lastModifiedColumn" :
  {
    "name"          : "LAST_UPDATED",
    "index"         : "empLastModIndexName"
  },
  "versionColumn"    :
  {
    "name"          : "VERSION_NUM",
    "method"        : "SEQUENTIAL"
  },
  "mediaTypeColumn"  :
  {
    "name"          : "CONTENT_TYPE"
  },
  "readOnly"         : true
}

```

See Also:

- [Key Assignment Method](#) (page 1-34)
 - [Versioning Method](#) (page 1-35)
 - *Oracle Database JSON Developer's Guide* for information about the syntax possibilities used by SQL condition `is json`
 - <http://tools.ietf.org/html/rfc4627> for the JSON RFC 4627 standard
 - *Oracle Database SecureFiles and Large Objects Developer's Guide* for information about SecureFiles LOB storage
-

1.5.1 Key Assignment Method

The key assignment method determines how keys are assigned to objects that are inserted into a collection.

Table 1-6 Key Assignment Methods

Method	Description
SEQUENCE	Keys are integers generated by a database sequence. You must specify the name of the sequence in the <code>keyColumn.sequenceName</code> field.
GUID	Keys are generated by the SQL function <code>SYS_GUID()</code> , which returns a globally unique RAW value (16 bytes). If necessary, the RAW value is converted to the SQL data type specified by <code>keyColumn.sqlType</code> .
UUID	Keys are generated by the built-in UUID capability of the Java Virtual Machine (JVM) on which the REST server is running, which returns a universally unique RAW value. If necessary, the RAW value is converted to the SQL data type specified by <code>keyColumn.sqlType</code> .
CLIENT	Keys are assigned by the client application (not recommended).

Oracle REST standards strongly recommend using server-assigned keys; that is, avoiding the key assignment method `CLIENT`. If you need simple numeric keys, Oracle recommends `SEQUENCE`. If any unique identifier is sufficient, Oracle recommends `UUID`.

If the key assignment method is `SEQUENCE`, `GUID`, or `UUID`, you insert a object into the collection with the operation [POST object](#) (page 1-26). The REST server always interprets `POST` as an insert operation, assigning a key and returning the key in the response body.

If the key assignment method is `CLIENT`, you cannot use `POST` to insert a object in the collection, because the URL path does not include the necessary key. Instead, you must insert the object into the collection using [PUT object](#) (page 1-24). If the object is not already in the collection, then the REST server interprets `PUT` as an insert operation. If the object is already in the collection, then the REST server interprets `PUT` as a replace operation. `PUT` is effectively equivalent to the SQL statement `MERGE`.

Caution:

If client-assigned keys are used and the key column type is `VARCHAR2` then Oracle recommends that the database character set be `AL32UTF8`. This ensures that conversion of the keys to the database character set is lossless.

Otherwise, if client-assigned keys contain characters that are not supported in your database character set then conversion of the key into the database character set during a read or write operation is lossy. This can lead to duplicate-key errors during insert operations. More generally, it can lead to unpredictable results. For example, a read operation could return a value that is associated with a different key from the one you expect.

1.5.2 Versioning Method

The versioning method determines how the REST server computes version values for objects when they are inserted into a collection or replaced.

Table 1-7 Versioning Methods

Method	Description
MD5	<p>The REST server computes an MD5 checksum on the bytes of object content. For bytes with character data types (such as <code>VARCHAR2</code> and <code>CLOB</code>), the computation uses UTF-8 encoding. For bytes with data type <code>BLOB</code>, the computation uses the encoding used to transmit the <code>POST</code> body, which can be either UTF-8 or UTF-16.</p> <p>For a bulk insert, the request body is parsed as an array of objects and the bytes of the individual objects are re-serialized with UTF-8 encoding, regardless of the encoding chosen for storage.</p> <p>In all cases, the checksum is computed on the bytes as they would be returned by a <code>GET</code> operation for the object.</p>

Table 1-7 (Cont.) Versioning Methods

Method	Description
SHA256 (default)	<p>The REST server computes a SHA256 checksum on the bytes of object content. For bytes with character data types (such as VARCHAR2 and CLOB), the computation uses UTF-8 encoding. For bytes with data type BLOB, the computation uses the encoding used to transmit the POST body, which can be either UTF-8 or UTF-16.</p> <p>For a bulk insert, the request body is parsed as an array of objects and the bytes of the individual objects are re-serialized with UTF-8 encoding, regardless of the encoding chosen for storage.</p> <p>In all cases, the checksum is computed on the bytes as they would be returned by a GET operation for the specific object.</p>
UUID	<p>Ignoring object content, the REST server generates a universally unique identifier (UUID)—a 32-character hexadecimal value—when the object is inserted and for every replace operation (even if the replace operation does not change the object content).</p>
TIMESTAMP	<p>Ignoring object content, the REST server generates an integer value, derived from the value returned by the SQL SYSTIMESTAMP function. The integer value changes at the level of accuracy of the system clock (typically microseconds or milliseconds).</p>
SEQUENTIAL	<p>Ignoring object content, the REST server assigns version 1 when the object is inserted and increments the version value every time the object is replaced.</p>
NONE	<p>The REST server does not assign version values during insert and replace operations. During GET operations, any non-null value stored in the version column is used as an ETag. Your application is responsible for populating the version column (using, for example, a PL/SQL trigger or asynchronous program).</p>

MD5 and SHA256 compute checksum values that change when the content itself changes, providing a very accurate way to invalidate client caches. However, they are costly, because the REST server must perform a byte-by-byte computation over the objects as they are inserted or replaced.

UUID is most efficient for input operations, because the REST server does not have to examine every byte of input or wait for SQL to return function values. However, replacement operations invalidate cached copies even if they do not change object content.

TIMESTAMP is useful when you need integer values or must compare two versions to determine which is more recent. As with UUID, replacement operations can invalidate cached copies without changing object content. Because the accuracy of the system clock may be limited, TIMESTAMP is not recommended if objects can change at very high frequency (many times per millisecond).

SEQUENTIAL is also useful when you need integer values or must compare two versions to determine which is more recent. Version values are easily understood by human users, and the version increases despite system clock limitations. However, the increment operation occurs within SQL; therefore, the new version value is not always available to be returned in the REST response body.

1.6 Security

ORDS, including SODA for REST, uses role-based access control, to secure services. The roles and privileges you need for SODA for REST are described here.

You should be familiar with the ORDS security features before reading this section. See *Oracle REST Data Services Installation, Configuration, and Development Guide* for the relevant information.

Database role **SODA_APP** must be granted to database users before they can use REST SODA. In addition, when a schema is enabled in ORDS using `ords.enable_schema`, a privilege is created such that only users with the application-server role `SODA Developer` can access the service. Specifically, `ords.enable_schema` creates the following privilege mapping:

```
exec ords.create_role('SODA Developer');
exec ords.create_privilege(p_name => 'oracle.soda.privilege.developer',
                          p_role_name => 'SODA Developer');
exec ords.create_privilege_mapping('oracle.soda.privilege.developer', '/soda/*');
```

This has the effect that, by default, a user must have the application-server role `SODA Developer` to access the JSON document store.

You can also add custom privilege mappings. For example:

```
declare
  l_patterns owa.vc_arr;
begin
  l_patterns(1) := '/soda/latest/employee';
  l_patterns(2) := '/soda/latest/employee/*';
  ords.create_role('EmployeeRole');
  ords.create_privilege(p_name      => 'EmployeePrivilege',
                       p_role_name => 'EmployeeRole');
  ords.create_privilege_mapping(p_privilege_name => 'EmployeePrivilege',
                                p_patterns      => l_patterns);

  commit;
end;
```

This example creates a privilege mapping that specifies that only users with role `EmployeeRole` can access the `employee` collection.

When multiple privilege patterns apply to the same resource, the privilege with the most specific pattern overrides the others. For example, patterns `/soda/latest/employees/*` and `/soda/*` both match the request URL, `http://example.org/ords/quine/soda/latest/employee/id1`.

Since `/soda/latest/employees/*` is more specific than `/soda/*`, only privilege `EmployeePrivilege` applies to the request.

Note:

`SODA_APP` is an Oracle Database role. `SODA Developer` is an application-server role.

1.6.1 Authentication Mechanisms

ORDS supports many different authentication mechanisms. JSON document store REST services are intended to be used in server-to-server interactions. Therefore, two-

legged OAuth (the client-credentials flow) is the recommended authentication mechanism to use with the JSON document store REST services. However, other mechanisms such as HTTP basic authentication, are also supported.

See Also:

Oracle REST Data Services Installation, Configuration, and Development Guide

1.6.2 Security Considerations for Development and Testing

Security considerations for development and testing are presented.

You can disable security and allow anonymous access by removing the default privilege mapping:

```
exec ords.delete_privilege_mapping('oracle.soda.privilege.developer', '/soda/*')
```

However, Oracle does *not* recommend that you allow anonymous access in production systems. That would allow an unauthenticated user to read, update, or drop any collection.

You can also use command `ords.war user` to create test users that have particular roles. For example (where *new_password* is a placeholder for the password for user bob):

```
# Create user bob with role SODA Developer
java -jar ords.war user bob "SODA Developer"
```

```
# Access the JSON document store as user bob using basic authentication
curl -u bob:new_password https://example.com/ords/scott/soda/latest/
```


C

collections

- creating, [1-6](#)
- deleting, [1-7](#)
- deleting documents from, [1-9](#)
- listing, [1-7](#)
- listing documents in, [1-11](#)
- specifications for, [1-30](#)

D

database role

- SODA_APP, [1-37](#)

DELETE collection operation, [1-24](#)

DELETE object operation, [1-25](#)

deleting collections, [1-7](#)

deleting documents from collections, [1-9](#)

documents

- deleting from collections, [1-9](#)
- filtering in collections, [1-14](#)
- inserting into collections
 - in bulk from JSON array, [1-10](#)
 - one at a time, [1-8](#)
- listing in collections, [1-11](#)
- retrieving from collections, [1-9](#)
- updating in collections, [1-13](#)

F

filtering documents in collections, [1-14](#)

G

GET collection operation, [1-19](#)

GET object operation, [1-22](#)

GET schema operation, [1-18](#)

I

inserting documents into collections

- in bulk from JSON array, [1-10](#)
- one at a time, [1-8](#)

installing SODA for REST, [1-3](#)

K

key assignment method, [1-34](#)

L

listing collections in schema, [1-7](#)

listing documents in collections, [1-11](#)

O

Oracle REST API, [1-1](#)

Oracle REST API HTTP operations

- response bodies, [1-17](#)
- summary of, [1-16](#)
- URI form for, [1-16](#)

P

POST array insert operation, [1-28](#)

POST bulk delete operation, [1-29](#)

POST object operation, [1-26](#)

POST query operation, [1-27](#)

PUT collection operation, [1-23](#)

PUT object operation, [1-24](#)

Q

query-by-example (QBE)

- examples, [1-14](#)

R

REST architectural style, [1-2](#)

retrieving documents from collections, [1-9](#)

S

security, [1-37](#)

SODA_APP database role, [1-37](#)

specifications
collection, [1-30](#)

U

updating documents in collections, [1-13](#)

V

versioning method, [1-35](#)