

OpenPaaS Database API Specification

REST VERSION

Contributors:

Rami Sellami, Telecom SudParis

Bruno Defude, Telecom SudParis

Table of content

1	OpenPaaS DataBase API : ODBAPI	3
1.1	Concepts comparison	3
1.2	Resources model of ODBAPI	4
1.3	Panorama of ODBAPI.....	4
1.4	Operations of ODBAPI.....	6
1.5	Detailed specification of ODBAPI	8
1.5.1	Get information about the user's access right	8
1.5.2	Get information about an environment	9
1.5.3	Get information about a database	9
1.5.4	Get information about an entity set	9
1.5.5	Get an entity set by its esName	10
1.5.6	Create an entity set	10
1.5.7	Delete an entity set	10
1.5.8	Get list of all entity sets	11
1.5.9	<i>Get an entity by its entityID</i>	12
1.5.10	<i>Update an entity</i>	12
1.5.11	<i>Create an entity</i>	13
1.5.12	Delete an entity	13
1.5.13	Get list of all entities.....	14

1 OpenPaaS DataBase API : ODBAPI

Nowadays, either relational DBMSs or NoSQL DBMSs are on demand and the application developer must be familiar with the proprietary API of each type of DBMS. However, these APIs are heterogeneous on different levels :

- API level: There are two categories of API. In the first category, the operations are not explicit in the API since the access queries to a database are described in strings (i.e. JDBC). In the second category, the operations are explained in the API with a method for each operation (i.e. MongoDB driver).
- Typing level: There are APIs that are closely related to a programming language and manipulate typed objects (i.e. Java JDBC) and others who handle more neutral data structures (i.e. JSON).

In order to satisfy different storage requirements, cloud applications usually need to access and interact with different relational and NoSQL data stores having heterogeneous APIs. This APIs heterogeneity induces two main problems. First it ties cloud applications to specific data stores hampering therefore their migration. Second, it requires developers to be familiar with different APIs.

For this sake, we propose OpenPaaS DataBase API (ODBAPI) a streamlined and a unified REST API enabling to execute CRUD operations on different NoSQL and relational databases. ODBAPI decouples cloud applications from data stores alleviating therefore their migration. Moreover it relieves developers task by removing the burden of managing different APIs. In the upcoming sections, we provide a specification and an implementation of ODBAPI based on a relational DBMS, a document data store called CouchDB, and a key/value data store called Riak.

1.1 Concepts comparison

In Table 1, we present an analogy between the concepts terminology in a relational DBMS, a key/value DBMS Riak, and a document DBMS CouchDB. Based on this, we propose our own concepts terminology in order to define the ODBAPI. For instance, a table in a relational database is equivalent to a database in a CouchDB and Riak databases. Hence, to eliminate this heterogeneity, we propose to rename this data structure Entity Set.

Table1 : Concepts comparison chart

Relational concepts	CouchDB concepts	Riak concepts	ODBAPI concepts
Database	Environment	Environment	Database
Table	Database	Database	Entity Set
Row	Document	Key/value	Entity
Column	Field	Value	Attribute

1.2 Resources model of ODBAPI

Based on Table 1, we propose the resources model defining the different target resources by ODBAPI (see Figure 1). The main entity in this model is the resource *environment* that is characterized by a name *envName*. An *environment* may contain one or multiple resources *database*. This latter is identified by a unique name *dbName* and contains one or more resources of type *entitySet*. An *entitySet* is determined by a unique name *esName* and has one or more resources of type *entity*. An *entity* is characterized by a unique *entityID* and has a set of resources of type *attribute*. This latter is identified by a content *attContent* that represents the value of an attribute.



Figure 1 ODBAPI resources model

1.3 Panorama of ODBAPI

In Figure 2, we present a panorama of ODBAPI. This API is designed to provide an abstraction layer and seamless interaction with data stores deployed in a cloud environment. Developers can execute CRUD operations in a uniform way regardless of whether the type of a data store is either relational or NoSQL. In Figure 2, we show mainly four parts that we introduce starting from the right side to the left side. In fact, we have first of all the deployed data stores (e.g. relational DBMS, Couch DB, etc.) that a developer may interact with. Second, we find the proprietary API and driver of each data store implemented by ODBAPI. For instance, we use in our API implementation the JDBC API and MySQL drive to interact with a relational DBMS. The third part of Figure 2 represents the ODBAPI implementation. In fact, this part represents the shared part between all the integrated data stores. In addition, it contains specific implementation of each data store. As we said previously, we propose in this paper a version implementing three data stores: (1) relational DBMS, (2) Couch DB, and (3) Riak by using their drivers and their appropriate APIs. How-

ever, to integrate a new data store and define interactions with it, one has simply to add the specific implementation of that data store. Finally, we show the different operations that ODBAPI offers to the user.

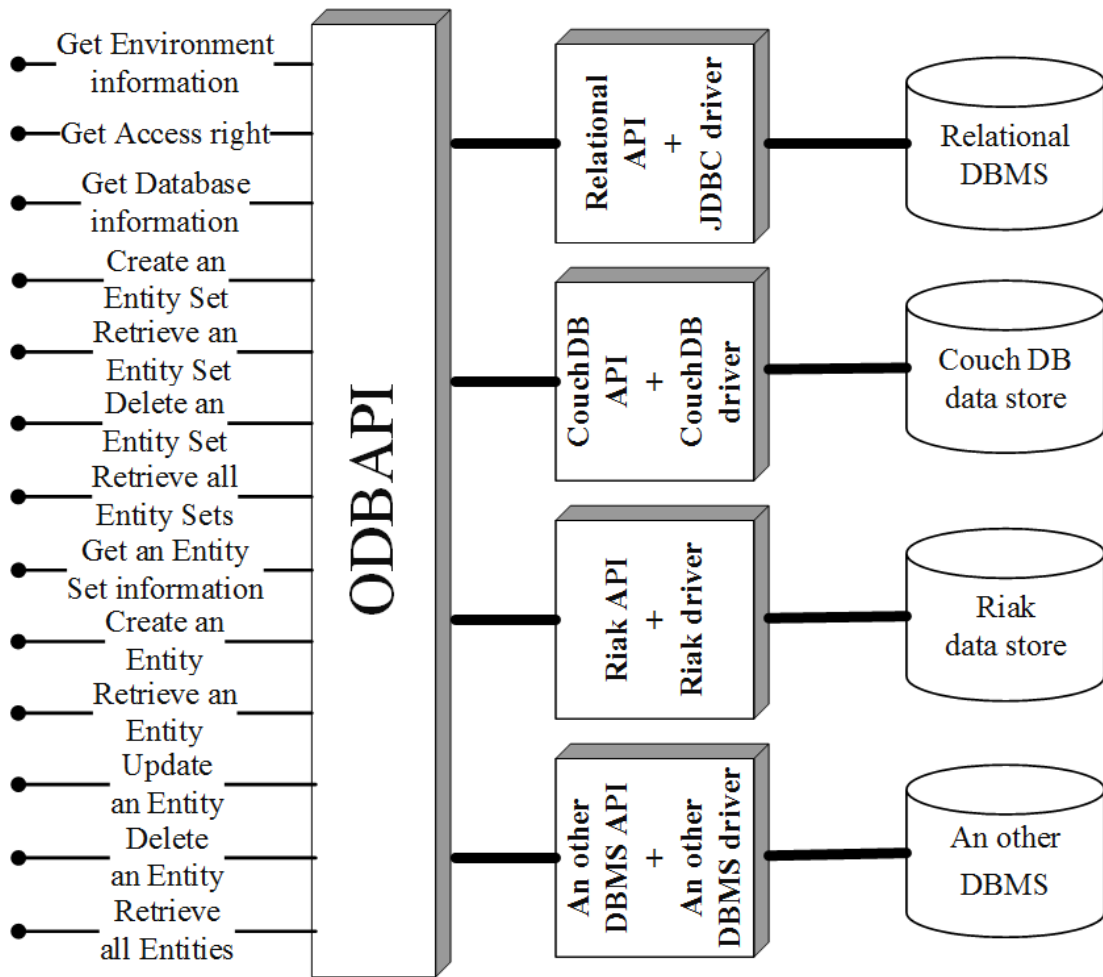


Figure 2 Overview of ODBAPI

In Figure3, we introduce the the protocol architecture of our API. In fact, the client interacts with cloud data stores through the ODBAPI based on REST protocol. Then, according to the target data store ODBAPI interacts with it based on the appropriate API. It is worthy noting that these proprietary APIs are not restful APIs.

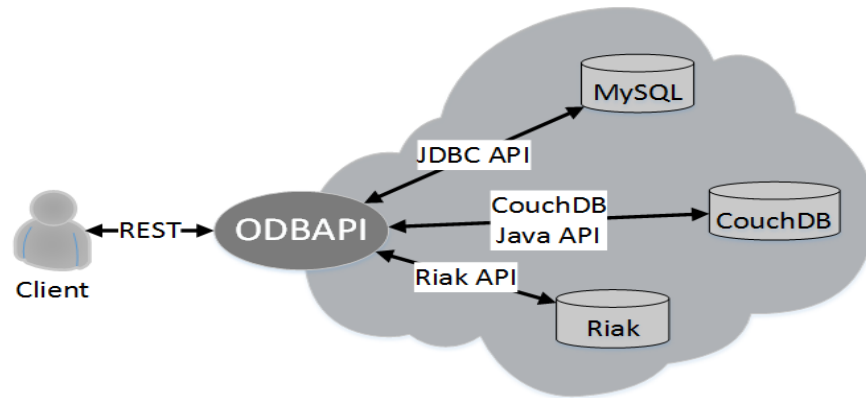


Figure 3 Protocol architecture of ODBAPI

1.4 Operations of ODBAPI

In Figure 4, we propose a box based representation of the different operations ensured by ODBAPI. Each box contains the name of a resource (e.g. /odbapi/{esName}, /odbapi/{esName}/{entityID}, etc) and the different operations that are intended to this resource. Each operation is ensured by a REST method (e.g. GET, PUT, etc).

In our specification, we consider two kinds of operations. The first operations family is dedicated to get meta-information about the resources using the REST method `GET`. In fact, ODBAPI offers four operations:

- Get information about the user's access right: This operation is provided by *getAccessRight* and allows a user to discover his access rights concerning the deployed data stores in a cloud environment. To do so, the user must append to his request the keyword *accessright*. This operation will help the user in choosing the appropriate data store according to his application requirements by listing his access right to databases. The *getAccessRight* operation is linked to the couple user/data store and an user should run it when he use data stores of an *environment* for the first time. In fact, once an user run it, its output is stored in order to avoid its rerun whenever the user interacts with the data stores. However, an user must re-run this operation whether there is some modifications in the data data stores.

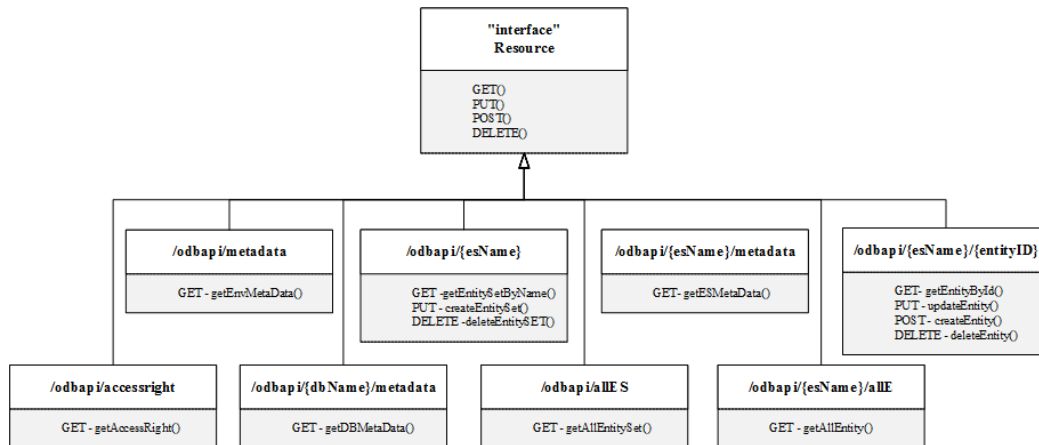


Figure 4 ODBAPI operations

- Get information about an *environment*: This operation is ensured by *getEnvMetadata* and lists the information about an *environment*. To execute this kind of operation, user must provide the keyword *metadata* in his request. This keyword should be also present in the following two operations. Following the execution of this operation, user discovers the deployed data stores in an *environment*. Thus, he will be able to choose the suitable data store. The *getEnvMetadata* operation is linked only to the different data stores in an *environment*.
- Get information about a *database*: An user can retrieve meta-information about a *database* by executing the operation *getDBMetadata* and providing the name of the target *database* *dbName*. This operation outputs information about a *database* (e.g. duplication, replication, etc) and the *entity sets* that contains.
- Get information about an *entity set*: This operation is provided by *getESMetadata* and enables to discover information about an *entity set* by giving its name *esName*. For instance, it helps the user to know the number of *entities* that an *entity set* contains and the *attributes* that constitute these *entities*.

Whereas the second operations family represents the CRUD operations executed on resources of type either *entitySet* or *Entity*. In this context, ODBAPI provides nine operations:

- Get an *entity set* by its *esName*: By executing the operation *getEntitySetByName*, an user can retrieve an *entity set* by giving its name *esName*. It is ensured by the *GET* method.
- Create an *entity set*: The operation *createEntitySet* allows an user to create an *entity set* by giving its name *esName*. This operation is provided by the REST method *PUT*.

- Delete an *entity set*: An *entity set* can be deleted by using the operation *deleteEntitySet* and giving as input its name *esName*. It is ensured by the *DELETE* method.
- Get list of all *entity sets*: User can retrieve the list of all *entity sets* by executing the operation *getAllEntitySet* and using the keyword *allES*. It outputs the names of the entity sets and several information (e.g. number of entities in each entity set, the type of database containing it, etc.).
- Get an *entity* by its *entityID*: By executing the operation *getEntityById*, an user can retrieve an *entity* by giving its identifier *entityID*. It is ensured by the *GET* method.
- Update an *entity*: An *entity* can be updated by using the operation *updateEntity* and its identifier *entityID*. It is ensured by the *PUT* method.
- Create an *entity*: The operation *createEntitySet* allows an user to create an *entity* by giving its identifier *entityID*. This operation is provided by the REST method *POST*.
- Delete an *entity*: An *entity* can be deleted by using the operation *deleteEntity* and giving as input its identifier *entityID*. It is ensured by the *DELETE* method.
- Get list of all *entities*: User can retrieve the list of all *entities* of an *entity set* by executing the operation *getAllEntity* and using the keyword *allE*. It outputs the identifiers of the *entities* and their contents.

1.5 Detailed specification of ODBAPI

1.5.1 Get information about the user's access right

Identifiant de la ressource	/odbapi/accessright/
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	Information about a user access right in JSON
Status code	200 if OK the error code otherwise

1.5.2 Get information about an environment

Identifiant de la ressource	/odbapi/metadata/
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	Metadata about an environment in JSON
Status code	200 if OK the error code otherwise

1.5.3 Get information about a database

Identifiant de la ressource	/odbapi/{db_name}/metadata
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	Metadata about a database in JSON
Status code	200 if OK the error code otherwise

1.5.4 Get information about an entity set

Identifiant de la ressource	/odbapi/{es_name}/metadata
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	Metadata about an entity set in JSON
Status code	200 if OK the error code otherwise

1.5.5 Create an entity set

Identifiant de la ressource	/odbapi/{es_name}/
HTTP method	PUT
Input parameter	Database-Type: Content-Type : Accept:
Status code	200 if OK the error code otherwise

```

Request
PUT /odbapi/person/

Database-Type: database/couchDB
Accept : application/json

Response
http/1.1 200 OK
Accept : application/json

The person entity set has been created successfully...

```

1.5.6 Get an entity set by its esName

Identifiant de la ressource	/odbapi/{es_name}/
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	Information about an entity set in JSON
Status code	200 if OK the error code otherwise

```

Request
GET /odbapi/person/

Database-Type: database/couchDB
Accept : application/json

Response
http/1.1 200 OK
Content-Type : application/json

{
  documentCount: 1,
  allDocuments: {
    total_rows: 1,
    offset: 0,
    rows: [1],

```

```

0: {
  id: "1",
  key: "1",
  value: {
    rev: "1-205b04ab704f7d5ff5e0fcc5f302fc41"
  }
}
updateSeq: 12,
name: "person"
}

```

1.5.7 Delete an entity set

Identifiant de la ressource	/odbapi/{es_name}/
HTTP method	DELETE
Input parameter	Database-Type: Content-Type : Accept:
Status code	200 if OK the error code otherwise

Request
DELETE /odbapi/person/

Database-Type: database/couchDB
Accept : application/json

Response
http/1.1 204 No Content
Accept : application/json

The person entity set has been deleted successfully...

1.5.8 Get list of all entity sets

Identifiant de la ressource	/odbapi/allES
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	List of all entity sets in JSON
Status code	200 if OK the error code otherwise

1.5.9 Create an entity

Identifiant de la ressource	/odbapi/{es_name}/{entityID}
HTTP method	POST
Input parameter	Database-Type: Content-Type : Accept: The content to insert in JSON
Status code	200 if OK the error code otherwise

```

Request
POST /odbapi/person/1

Database-Type: database/couchDB
Content-Type: application/json
Accept: application/json

{
  "name": "john",
  "surname": "Doe"
}

Response
http/1.1 200 OK
Accept : application/json

Document 1 has been created successfully...

```

1.5.10 Get an entity by its entityID

Identifiant de la ressource	/odbapi/{es_name}/{entityID}
HTTP method	GET
Input parameter	Database-Type: Accept:
Response	The content of an entity in JSON
Status code	200 if OK the error code otherwise

```

Request
GET /odbapi/person/1

Database-Type: database/couchDB
Content-Type: application/json
Accept: application/json

Response

```

```

http/1.1 200 OK
Accept : application/json

{
  _id: "1",
  _rev: "1-205b04ab704f7d5ff5e0fcc5f302fc41",
  name: "john",
  surname: "Doe"
}

```

1.5.11 Update an entity

Identifiant de la ressource	/odbapi/{es_name}/{entityID}
HTTP method	PUT
Input parameter	Database-Type: Content-Type : Accept:
Status code	200 if OK the error code otherwise

```

Request
PUT /odbapi/person/1

Database-Type: database/couchDB
Content-Type: application/json
Accept: application/json

{
  "name": "Peter",
  "surname": "Doe"
}

Response
http/1.1 200 OK
Accept : application/json

Document 1 has been updated successfully...

```

1.5.12 Delete an entity

Identifiant de la ressource	/odbapi/{es_name}/{entityID}
HTTP method	DELETE
Input parameter	Database-Type: Accept:
Status code	200 if OK the error code otherwise

```

Request
DELETE /odbapi/person/1

Database-Type: database/couchDB
Accept: application/json

Response
http/1.1 200 OK
Accept : application/json

Document 1 has been deleted successfully...

```

1.5.13 Get list of all entities

Identifiant de la ressource	/odbapi/{es_name}/allE
HTTP method	GET
Input parameter	Database-Type: Content-Type : Accept:
Response	List of all entities in JSON
Status code	200 if OK the error code otherwise

1.6 ODBAPI use cases

1.6.1 First scenario: Application migration from one data store to another

In Figure 5, we exemplify a migration scenario where the Application A needs to migrate from one cloud environment where it interacts with the document data store CouchDB to another in order to meet new data requirements. In the new cloud environment, the application connects to another document data store Mongo DB.

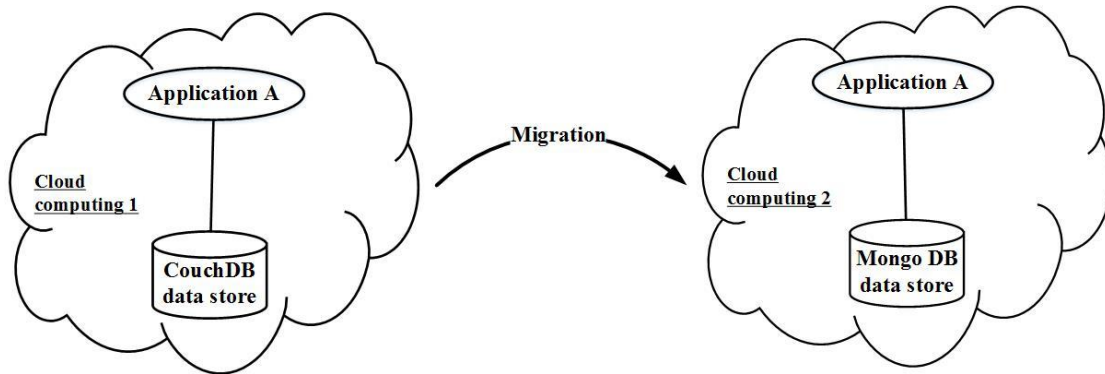


Figure 5 Application migration from on cloud environment to another scenario

In the source code below, we present an example of using ODBAPI in the case of migration from couchDB data store to MongoDB data store. Indeed, the developer creates two entities in a couchDB data store. Then after the migration, developer still using ODBAPI and he has just to replace the old type of the target data store (i.e. *database/couchDB*) by the new target data store (i.e. *database/mangoDB*). Hence, ODBAPI simplifies application migration. In fact, it decouples cloud applications from data stores alleviating therefore their migration.

```

/*****
 Create 2 documents in a couchDB database
 *****/
CreateEntitySetImpl ces = new CreateEntitySetImpl();

/*****Insert document 1 in couchDB*****/
InputStream is = new FileInputStream("file/EntitySet1.json");
InputStreamReader isr =new InputStreamReader(is);
JSONTokener tokener = new JSONTokener(isr);
JSONObject jsonEntity = new JSONObject(tokener);
ces.createEntitySet("http://localhost:8182/odbapi/person/1","database/couchDB", jsonEntity);

/*****Insert document 2 in couchDB*****/
InputStream is1 = new FileInputStream("file/EntitySet2.json");
InputStreamReader isr1 =new InputStreamReader(is1);
JSONTokener tokener1 = new JSONTokener(isr1);
JSONObject jsonEntity1 = new JSONObject(tokener1);
ces.createEntitySet("http://localhost:8182/odbapi/person/2","database/couchDB", jsonEntity1);

/*****
-----<Migration from Couchdb to MangoDB>-----
 Retrieve 2 documents in a MangoDB database

 *****/
RetrieveEntityImpl re = new RetrieveEntityImpl();

/*****Retrieve document 1 in couchDB*****/
re.retrieveEntity("http://localhost:8182/odbapi/person/1","database/mangoDB");

/*****Retrieve document 1 in couchDB*****/
re.retrieveEntity("http://localhost:8182/odbapi/person/2","database/mangoDB");

```

1.6.2 Second scenario: Polyglot persistence

In a cloud environment, an application can use multiple data stores that correspond to what is popularly referred to as the polyglot persistence. In Figure 6, we show an example of this situation. Application A interacts with three heterogeneous data stores: a relational data store, a document data store that is CouchDB, and a key value data store which is Riak.




Figure 6 Multiple data store use in one cloud environment

In the source code below, we present an example of using ODBAPI in the case of interacting of two data stores at the same time. These data stores are CouchDB and MySQL.

```

/*****
Polyglot persistence
*****/
CreateEntitySetImpl ces = new CreateEntitySetImpl();

/*****Insert document 1 in couchDB database*****/
InputStream is = new FileInputStream("file/EntitySet1.json");
InputStreamReader isr =new InputStreamReader(is);
JSONTokener tokener = new JSONTokener(isr);
JSONObject jsonEntity = new JSONObject(tokener);
ces.createEntitySet("http://localhost:8182/odbapi/person/1", "database/couchDB", jsonEntity);

/*****Retrieve document 1 in couchDB database*****/
RetrieveEntityImpl re = new RetrieveEntityImpl();
re.retrieveEntity("http://localhost:8182/odbapi/person/1", "database/ couchDB");

/*****Insert a tuple in MySQL database*****/
InputStream is1 = new FileInputStream("file/EntitySet1.json");
InputStreamReader isr1 =new InputStreamReader(is1);
JSONTokener tokener1 = new JSONTokener(isr1);
JSONObject jsonEntity1 = new JSONObject(tokener1);
ces.createEntitySet("http://localhost:8182/odbapi/city/2 ", "database/MySQL", "world",
jsonEntity1);

/*****Retrieve a tuple in MySQL database *****/
re.retrieveEntity("http://localhost:8182/odbapi/city/2 ", "database/MySQL", "world");
  
```