

# Quarkus - Reactive SQL Clients

The Reactive SQL Clients have a straightforward API focusing on scalability and low-overhead. Currently, the following database servers are supported:

- DB2
- PostgreSQL
- MariaDB/MySQL

In this guide, you will learn how to implement a simple CRUD application exposing data stored in **PostgreSQL** over a RESTful API.



Extension and connection pool class names for each client can be found at the bottom of this document.



If you are not familiar with the Quarkus Vert.x extension, consider reading the [Using Eclipse Vert.x](#) guide first.

The application shall manage fruit entities:

```
public class Fruit {  
    public Long id;  
    public String name;  
    public Fruit() {  
    }  
    public Fruit(String name) {  
        this.name = name;  
    }  
    public Fruit(Long id, String name) {  
        this.id = id;  
        this.name = name;  
    }  
}
```



Do you need a ready-to-use PostgreSQL server to try out the examples?

```
docker run --ulimit memlock=-1:-1 -it --rm=true
--memory-swappiness=0 --name quarkus_test -e
POSTGRES_USER=quarkus_test -e
POSTGRES_PASSWORD=quarkus_test -e
POSTGRES_DB=quarkus_test -p 5432:5432 postgres:10.5
```

## Installing

### Reactive PostgreSQL Client extension

First, make sure your project has the `quarkus-reactive-pg-client` extension enabled. If you are creating a new project, set the `extensions` parameter as follows:

```
mvn io.quarkus:quarkus-maven-plugin:1.7.3.Final:create \
  -DprojectId=org.acme \
  -DprojectArtifactId=reactive-pg-client-quickstart \
  -Dextensions="reactive-pg-client"
cd reactive-pg-client-quickstart
```

If you have an already created project, the `reactive-pg-client` extension can be added to an existing Quarkus project with the `add-extension` command:

```
./mvnw quarkus:add-extension -Dextensions="reactive-pg-client"
```

Otherwise, you can manually add this to the dependencies section of your `pom.xml` file:

```
<dependency>
  <groupId>io.quarkus</groupId>
  <artifactId>quarkus-reactive-pg-client</artifactId>
</dependency>
```

### Mutiny

Reactive REST endpoints in your application that return Uni or Multi need `Mutiny support for RESTEasy` extension (`io.quarkus:quarkus-resteasy-mutiny`) to work properly:

```
./mvnw quarkus:add-extension -Dextensions="resteasy-mutiny"
```



In this guide, we will use the Mutiny API of the Reactive PostgreSQL Client. If you're not familiar with Mutiny reactive types, read the [Getting Started with Reactive guide](#) first.

## JSON Binding

We will expose `Fruit` instances over HTTP in the JSON format. Consequently, you also need to add the `quarkus-resteasy-jsonb` extension:

```
./mvnw quarkus:add-extension -Dextensions="resteasy-jsonb"
```

If you prefer not to use the command line, manually add this to the dependencies section of your `pom.xml` file:

```
<dependency>
  <groupId>io.quarkus</groupId>
  <artifactId>quarkus-resteasy-jsonb</artifactId>
</dependency>
```

Of course, this is only a requirement for this guide, not any application using the Reactive PostgreSQL Client.

## Configuring

The Reactive PostgreSQL Client can be configured with standard Quarkus datasource properties and a reactive URL:

*src/main/resources/application.properties*

```
quarkus.datasource.db-kind=postgresql
quarkus.datasource.username=quarkus_test
quarkus.datasource.password=quarkus_test
quarkus.datasource.reactive.url=postgresql://localhost:5432/quarkus_test
```

With that you may create your `FruitResource` skeleton and `@Inject` a `io.vertx.mutiny.pgclient.PgPool` instance:

src/main/java/org/acme/vertx/FruitResource.java

```
@Path("fruits")
@Produces(MediaType.APPLICATION_JSON)
@Consumes(MediaType.APPLICATION_JSON)
public class FruitResource {

    @Inject
    io.vertx.mutiny.pgclient.PgPool client;
}
```

## Database schema and seed data

Before we implement the REST endpoint and data management code, we need to setup the database schema. It would also be convenient to have some data inserted upfront.

For production we would recommend to use something like the [Flyway database migration tool](#). But for development we can simply drop and create the tables on startup, and then insert a few fruits.

src/main/java/org/acme/vertx/FruitResource.java

```
@Inject
@ConfigProperty(name = "myapp.schema.create", defaultValue =
"true") ①
boolean schemaCreate;

@PostConstruct
void config() {
    if (schemaCreate) {
        initdb();
    }
}

private void initdb() {
    // TODO
}
```



You may override the default value of the `myapp.schema.create` property in the `application.properties` file.

Almost ready! To initialize the DB in development mode, we will use the client simple `query` method. It returns a `Uni` and thus can be composed to execute queries sequentially:

```

client.query("DROP TABLE IF EXISTS fruits").execute()
    .flatMap(r -> client.query("CREATE TABLE fruits (id SERIAL
PRIMARY KEY, name TEXT NOT NULL)").execute())
    .flatMap(r -> client.query("INSERT INTO fruits (name) VALUES
('Orange')").execute())
    .flatMap(r -> client.query("INSERT INTO fruits (name) VALUES
('Pear')").execute())
    .flatMap(r -> client.query("INSERT INTO fruits (name) VALUES
('Apple')").execute())
    .await().indefinitely();

```



#### *Breaking Change in Quarkus 1.5*

Vert.x 3.9, integrated in Quarkus 1.5, introduces a breaking change regarding the **query** methods. To retrieve the result you need to call **.execute()**.



Wondering why we need block until the latest query is completed? This code is part of a **@PostConstruct** method and Quarkus invokes it synchronously. As a consequence, returning prematurely could lead to serving requests while the database is not ready yet.

That's it! So far we have seen how to configure a pooled client and execute simple queries. We are now ready to develop the data management code and implement our RESTful endpoint.

## Using

### Query results traversal

In development mode, the database is set up with a few rows in the **fruits** table. To retrieve all the data, we will use the **query** method again:

```

Uni<RowSet<Row>> rowSet = client.query("SELECT id, name FROM fruits
ORDER BY name ASC").execute();

```

When the operation completes, we will get a **RowSet** that has all the rows buffered in memory. A **RowSet** is an **java.lang.Iterable<Row>** and thus can be converted to a **Multi**:

```

Multi<Fruit> fruits = rowSet
    .onItem().transformToMulti(set ->
Multi.createFrom().iterable(set))
    .onItem().transform(Fruit::from);

```

The **Fruit#from** method converts a **Row** instance to a **Fruit** instance. It is extracted as a convenience for the implementation of the other data management methods:

*src/main/java/org/acme/vertx/Fruit.java*

```
private static Fruit from(Row row) {  
    return new Fruit(row.getLong("id"), row.getString("name"));  
}
```

Putting it all together, the `Fruit.findAll` method looks like:

*src/main/java/org/acme/vertx/Fruit.java*

```
public static Multi<Fruit> findAll(PgPool client) {  
    return client.query("SELECT id, name FROM fruits ORDER BY name  
ASC").execute()  
        .onItem().transformToMulti(set ->  
Multi.createFrom().iterable(set))  
        .onItem().transform(Fruit::from);  
}
```

And the endpoint to get all fruits from the backend:

*src/main/java/org/acme/vertx/FruitResource.java*

```
@GET  
public Multi<Fruit> get() {  
    return Fruit.findAll(client);  
}
```

Now start Quarkus in `dev` mode with:

```
./mvnw compile quarkus:dev
```

Lastly, open your browser and navigate to <http://localhost:8080/fruits>, you should see:

```
[{"id":3,"name":"Apple"}, {"id":1,"name":"Orange"}, {"id":2,"name":"Pear"}]
```

## Prepared queries

The Reactive PostgreSQL Client can also prepare queries and take parameters that are replaced in the SQL statement at execution time:

```
client.preparedQuery("SELECT id, name FROM fruits WHERE id =  
$1").execute(Tuple.of(id))
```



The SQL string can refer to parameters by position, using \$1, \$2, ...etc.

Similar to the simple `query` method, `preparedQuery` returns an instance of `PreparedQuery<RowSet<Row>>`. Equipped with this tooling, we are able to safely use an `id` provided by the user to get the details of a particular fruit:

*src/main/java/org/acme/vertx/Fruit.java*

```
public static Uni<Fruit> findById(PgPool client, Long id) {  
    return client.preparedQuery("SELECT id, name FROM fruits WHERE  
id = $1").execute(Tuple.of(id)) ①  
        .onItem().transform(RowSet::iterator) ②  
        .onItem().transform(iterator -> iterator.hasNext() ?  
from(iterator.next()) : null); ③  
}
```

- ① Create a `Tuple` to hold the prepared query parameters.
- ② Get an `Iterator` for the `RowSet` result.
- ③ Create a `Fruit` instance from the `Row` if an entity was found.

And in the JAX-RS resource:

*src/main/java/org/acme/vertx/FruitResource.java*

```
@GET  
@Path("{id}")  
public Uni<Response> getSingle(@PathParam Long id) {  
    return Fruit.findById(client, id)  
        .onItem().transform(fruit -> fruit != null ?  
Response.ok(fruit) : Response.status(Status.NOT_FOUND)) ①  
        .onItem().transform(ResponseBuilder::build); ②  
}
```

- ① Prepare a JAX-RS response with either the `Fruit` instance if found or the `404` status code.
- ② Build and send the response.

The same logic applies when saving a `Fruit`:

*src/main/java/org/acme/vertx/Fruit.java*

```
public Uni<Long> save(PgPool client) {  
    return client.preparedQuery("INSERT INTO fruits (name) VALUES  
($1) RETURNING (id)").execute(Tuple.of(name))  
        .onItem().transform(pgRowSet ->  
pgRowSet.iterator().next().getLong("id"));  
}
```

And in the web resource we handle the **POST** request:

*src/main/java/org/acme/vertx/FruitResource.java*

```
@POST
public Uni<Response> create(Fruit fruit) {
    return fruit.save(client)
        .onItem().transform(id -> URI.create("/fruits/" + id))
        .onItem().transform(uri ->
Response.created(uri).build());
}
```

## Result metadata

A **RowSet** does not only hold your data in memory, it also gives you some information about the data itself, such as:

- the number of rows affected by the query (inserted/deleted/updated/retrieved depending on the query type),
- the column names.

Let's use this to support removal of fruits in the database:

*src/main/java/org/acme/vertx/Fruit.java*

```
public static Uni<Boolean> delete(PgPool client, Long id) {
    return client.preparedQuery("DELETE FROM fruits WHERE id =
$1").execute(Tuple.of(id))
        .onItem().transform(pgRowSet -> pgRowSet.rowCount() ==
1); ❶
}
```

❶ Inspect metadata to determine if a fruit has been actually deleted.

And to handle the HTTP **DELETE** method in the web resource:

*src/main/java/org/acme/vertx/FruitResource.java*

```
@DELETE
@Path("/{id}")
public Uni<Response> delete(@PathParam Long id) {
    return Fruit.delete(client, id)
        .onItem().transform(deleted -> deleted ?
Status.NO_CONTENT : Status.NOT_FOUND)
        .onItem().transform(status ->
Response.status(status).build());
}
```



With **GET**, **POST** and **DELETE** methods implemented, we may now create a minimal web page to try the RESTful application out. We will use **jQuery** to simplify interactions with the backend:

```
<!doctype html>
<html>
<head>
  <meta charset="utf-8"/>
  <title>Reactive PostgreSQL Client - Quarkus</title>
  <script src="https://code.jquery.com/jquery-3.3.1.min.js"
    integrity="sha256-
FgpCb/KJQlLnF0u91ta32o/NMZxltwRo8QtmkMRdAu8="
crossorigin="anonymous"></script>
  <script type="application/javascript" src="fruits.js"></script>
</head>
<body>

<h1>Fruits API Testing</h1>

<h2>All fruits</h2>
<div id="all-fruits"></div>

<h2>Create Fruit</h2>
<input id="fruit-name" type="text">
<button id="create-fruit-button" type="button">Create</button>
<div id="create-fruit"></div>

</body>
</html>
```

In the Javascript code, we need a function to refresh the list of fruits when:

- the page is loaded, or
- a fruit is added, or
- a fruit is deleted.

```

function refresh() {
    $.get('/fruits', function (fruits) {
        var list = '';
        (fruits || []).forEach(function (fruit) { ❶
            list = list
                + '<tr>'
                + '<td>' + fruit.id + '</td>'
                + '<td>' + fruit.name + '</td>'
                + '<td><a href="#" onclick="deleteFruit(' +
fruit.id + ')">Delete</a></td>'
                + '</tr>'
        });
        if (list.length > 0) {
            list = ''
                +
'<table><thead><th>Id</th><th>Name</th><th></th></thead>'
                + list
                + '</table>';
        } else {
            list = "No fruits in database"
        }
        $('#all-fruits').html(list);
    });
}

function deleteFruit(id) {
    $.ajax('/fruits/' + id, {method: 'DELETE'}).then(refresh);
}

$(document).ready(function () {

    $('#create-fruit-button').click(function () {
        var fruitName = $('#fruit-name').val();
        $.post({
            url: '/fruits',
            contentType: 'application/json',
            data: JSON.stringify({name: fruitName})
        }).then(refresh);
    });

    refresh();
});

```

❶ The **fruits** parameter is not defined when the database is empty.

All done! Navigate to <http://localhost:8080/fruits.html> and read/create/delete some fruits.

# Database Clients details

Database	Extension name	Pool class name
DB2	<code>quarkus-reactive-db2-client</code>	<code>io.vertx.mutiny.db2client.DB2Pool</code>
MariaDB/MySQL	<code>quarkus-reactive-mysql-client</code>	<code>io.vertx.mutiny.mysqlclient.MySQLPool</code>
PostgreSQL	<code>quarkus-reactive-pg-client</code>	<code>io.vertx.mutiny.pgclient.PgPool</code>

## Transactions

The reactive SQL clients support transactions. A transaction is started with `client.begin()` and terminated with either `tx.commit()` or `tx.rollback()`. All these operations are asynchronous:

- `client.begin()` returns a `Uni<Transaction>`
- `client.commit()` and `client.rollback()` return `Uni<Void>`

Managing transactions in the reactive programming world can be cumbersome. Instead of writing repetitive and complex (thus error-prone!) code, you can use the `io.vertx.mutiny.sqlclient.SqlClientHelper`.

The following snippet shows how to run 2 insertions in the same transaction:

```
public static Uni<Void> insertTwoFruits(PgPool client, Fruit
fruit1, Fruit fruit2) {
    return SqlClientHelper.inTransactionUni(client, tx -> {
        Uni<RowSet<Row>> insertOne = tx.preparedQuery("INSERT INTO
fruits (name) VALUES ($1) RETURNING (id)")
            .execute(Tuple.of(fruit1.name));
        Uni<RowSet<Row>> insertTwo = tx.preparedQuery("INSERT INTO
fruits (name) VALUES ($1) RETURNING (id)")
            .execute(Tuple.of(fruit2.name));

        return insertOne.and(insertTwo)
            // Ignore the results (the two ids)
            .onItem().ignore().andContinueWithNull();
    });
}
```

In this example, the transaction is automatically committed on success or rolled back on failure.

You can also create dependent actions as follows:

```

return SqlClientHelper.inTransactionUni(client, tx -> tx
    .preparedQuery("INSERT INTO person (firstname,lastname)
VALUES ($1,$2) RETURNING (id)")
    .execute(Tuple.of(person.getFirstName(),
person.getLastName())))

    .onItem().transformToUni(id -> tx.preparedQuery("INSERT
INTO addr (person_id,addrline1) VALUES ($1,$2)")

    .execute(Tuple.of(id.iterator().next().getLong("id"),
person.getLastName())))

    .onItem().ignore().andContinueWithNull());

```

## UNIX Domain Socket connections

The PostgreSQL and MariaDB/MySQL clients can be configured to connect to the server through a UNIX domain socket.

First make sure that [native transport support](#) is enabled.

Then configure the database connection url. This step depends on the database type.

### PostgreSQL

PostgreSQL domain socket paths have the following form: `<directory>/.s.PGSQL.<port>`

The database connection url must be configured so that:

- the `host` is the `directory` in the socket path
- the `port` is the `port` in the socket path

Consider the following socket path: `/var/run/postgresql/.s.PGSQL.5432`.

In `application.properties` add:

```

quarkus.datasource.reactive.url=postgresql://:5432/quarkus_test?hos
t=/var/run/postgresql

```

### MariaDB/MySQL

The database connection url must be configured so that the `host` is the socket path.


Consider the following socket path: `/var/run/mysqld/mysqld.sock`.




In `application.properties` add:


```
quarkus.datasource.reactive.url=mysql:///quarkus_test?host=/var/run
/mysql/mysql.sock
```

# Configuration Reference


## Common Datasource


 Configuration property fixed at build time - All other configuration properties are overridable at runtime

Configuration property	Type	Default
 <code>quarkus.datasource.db-kind</code>  The kind of database we will connect to (e.g. h2, postgresql...).	string	
 <code>quarkus.datasource.health.enabled</code>  Whether or not an health check is published in case the smallrye-health extension is present. This is a global setting and is not specific to a datasource.	boolean	<code>true</code>
 <code>quarkus.datasource.metrics.enabled</code>  Whether or not datasource metrics are published in case the smallrye-metrics extension is present. This is a global setting and is not specific to a datasource. NOTE: This is different from the "jdbc.enable-metrics" property that needs to be set on the JDBC datasource level to enable collection of metrics for that datasource.	boolean	<code>false</code>
<code>quarkus.datasource.max-size</code>	int	<code>20</code>
<code>quarkus.datasource.username</code>  The datasource username	string	
<code>quarkus.datasource.password</code>  The datasource password	string	
<code>quarkus.datasource.credentials-provider</code>  The credentials provider name	string	

<code>quarkus.datasource.credentials-provider-name</code>		
The credentials provider bean name. It is the <code>&amp;#64;Named</code> value of the credentials provider bean. It is used to discriminate if multiple CredentialsProvider beans are available. For Vault it is: vault-credentials-provider. Not necessary if there is only one credentials provider available.	string	
<b>Additional named datasources</b>	<b>Type</b>	<b>Default</b>
 <code>quarkus.datasource."datasource-name".db-kind</code>		
The kind of database we will connect to (e.g. h2, postgresql...).	string	
<code>quarkus.datasource."datasource-name".max-size</code>	int	20
<code>quarkus.datasource."datasource-name".username</code>		
The datasource username	string	
<code>quarkus.datasource."datasource-name".password</code>		
The datasource password	string	
<code>quarkus.datasource."datasource-name".credentials-provider</code>		
The credentials provider name	string	
<code>quarkus.datasource."datasource-name".credentials-provider-name</code>		
The credentials provider bean name. It is the <code>&amp;#64;Named</code> value of the credentials provider bean. It is used to discriminate if multiple CredentialsProvider beans are available. For Vault it is: vault-credentials-provider. Not necessary if there is only one credentials provider available.	string	

## Reactive Datasource

 Configuration property fixed at build time - All other configuration properties are overridable at runtime

<b>Configuration property</b>	<b>Type</b>	<b>Default</b>
 <code>quarkus.datasource.reactive</code>		
If we create a Reactive datasource for this datasource.	boolean	true

<code>quarkus.datasource.reactive.cache-prepared-statements</code>	boolean	<code>false</code>
Whether prepared statements should be cached on the client side.		
<code>quarkus.datasource.reactive.url</code>	string	
The datasource URL.		
<code>quarkus.datasource.reactive.max-size</code>	int	
The datasource pool maximum size.		
<code>quarkus.datasource.reactive.trust-all</code>	boolean	<code>false</code>
Whether all server certificates should be trusted.		
<code>quarkus.datasource.reactive.trust-certificate-pem</code>	boolean	<code>false</code>
PEM Trust config is disabled by default.		
<code>quarkus.datasource.reactive.trust-certificate-pem.certs</code>	list of string	
Comma-separated list of the trust certificate files (Pem format).		
<code>quarkus.datasource.reactive.trust-certificate-jks</code>	boolean	<code>false</code>
JKS config is disabled by default.		
<code>quarkus.datasource.reactive.trust-certificate-jks.path</code>	string	
Path of the key file (JKS format).		
<code>quarkus.datasource.reactive.trust-certificate-jks.password</code>	string	
Password of the key file.		
<code>quarkus.datasource.reactive.trust-certificate-pfx</code>	boolean	<code>false</code>
PFX config is disabled by default.		
<code>quarkus.datasource.reactive.trust-certificate-pfx.path</code>	string	
Path to the key file (PFX format).		
<code>quarkus.datasource.reactive.trust-certificate-pfx.password</code>	string	
Password of the key.		

<code>quarkus.datasource.reactive.key-certificate-pem</code> PEM Key/cert config is disabled by default.	boolean	<code>false</code>
<code>quarkus.datasource.reactive.key-certificate-pem.keys</code> Comma-separated list of the path to the key files (Pem format).	list of string	
<code>quarkus.datasource.reactive.key-certificate-pem.certs</code> Comma-separated list of the path to the certificate files (Pem format).	list of string	
<code>quarkus.datasource.reactive.key-certificate-jks</code> JKS config is disabled by default.	boolean	<code>false</code>
<code>quarkus.datasource.reactive.key-certificate-jks.path</code> Path of the key file (JKS format).	string	
<code>quarkus.datasource.reactive.key-certificate-jks.password</code> Password of the key file.	string	
<code>quarkus.datasource.reactive.key-certificate-pfx</code> PFX config is disabled by default.	boolean	<code>false</code>
<code>quarkus.datasource.reactive.key-certificate-pfx.path</code> Path to the key file (PFX format).	string	
<code>quarkus.datasource.reactive.key-certificate-pfx.password</code> Password of the key.	string	
<code>quarkus.datasource.reactive.thread-local</code> Experimental: use one connection pool per thread.	boolean	
<code>quarkus.datasource.reactive.reconnect-attempts</code> The number of reconnection attempts when a pooled connection cannot be established on first try.	int	<code>0</code>



<code>quarkus.datasource.reactive.reconnect-interval</code>	Duration ?	PT1S
The interval between reconnection attempts when a pooled connection cannot be established on first try.		
<code>quarkus.datasource.reactive.idle-timeout</code>	Duration ?	no timeou t
The maximum time without data written to or read from a connection before it is removed from the pool.		



#### About the Duration format

The format for durations uses the standard `java.time.Duration` format. You can learn more about it in the [Duration#parse\(\) javadoc](#).

You can also provide duration values starting with a number. In this case, if the value consists only of a number, the converter treats the value as seconds. Otherwise, `PT` is implicitly prepended to the value to obtain a standard `java.time.Duration` format.

## DB2

Configuration property fixed at build time - All other configuration properties are overridable at runtime

Configuration property	Type	Default
<code>quarkus.datasource.reactive.db2.ssl</code>	boolean	false
Whether SSL/TLS is enabled.		


## MariaDB/MySQL

Configuration property fixed at build time - All other configuration properties are overridable at runtime

Configuration property	Type	Default
<code>quarkus.datasource.reactive.mysql.charset</code>	string	
Charset for connections.		
<code>quarkus.datasource.reactive.mysql.collation</code>	string	
Collation for connections.		

<code>quarkus.datasource.reactive.mysql.ssl-mode</code> Desired security state of the connection to the server. See <a href="#">MySQL Reference Manual</a> .	disabled, preferred, required, verify- -ca, verify- -identity	disabled
---	---	----------

## PostgreSQL

 Configuration property fixed at build time - All other configuration properties are overridable at runtime

Configuration property	Type	Default
<code>quarkus.datasource.reactive.postgresql.pipelining-limit</code> The maximum number of inflight database commands that can be pipelined.	int	
<code>quarkus.datasource.reactive.postgresql.ssl-mode</code> SSL operating mode of the client. See <a href="#">Protection Provided in Different Modes</a> .	disabled, allow, preferred, required, verify- -ca, verify- -full	disabled