

Quarkus - Scheduling Periodic Tasks with Quartz

Modern applications often need to run specific tasks periodically. In this guide, you learn how to schedule periodic clustered tasks using the [Quartz](#) extension.



This technology is considered preview.

In *preview*, backward compatibility and presence in the ecosystem is not guaranteed. Specific improvements might require to change configuration or APIs and plans to become *stable* are under way. Feedback is welcome on our [mailing list](#) or as issues in our [GitHub issue tracker](#).

For a full list of possible extension statuses, check our [FAQ entry](#).



If you only need to run in-memory scheduler use the [Scheduler](#) extension.

Prerequisites

To complete this guide, you need:

- less than 10 minutes
- an IDE
- JDK 1.8+ installed with `JAVA_HOME` configured appropriately
- Apache Maven 3.5.3+
- Docker and Docker Compose installed on your machine

Architecture

In this guide, we are going to expose one Rest API `tasks` to visualise the list of tasks created by a Quartz job running every 10 seconds.

Solution

We recommend that you follow the instructions in the next sections and create the application step by step. However, you can go right to the completed example.

Clone the Git repository: `git clone https://github.com/quarkusio/quarkus-quickstarts.git`, or download an [archive](#).

The solution is located in the `quartz-quickstart` directory.

Creating the Maven project

First, we need a new project. Create a new project with the following command:

```
mvn io.quarkus:quarkus-maven-plugin:1.2.0.CR1:create \
  -DprojectId=org.acme \
  -DprojectArtifactId=quartz-quickstart \
  -DclassName="org.acme.quartz.TaskResource" \
  -Dpath="/tasks" \
  -Dextensions="quartz, hibernate-orm-panache, flyway, resteasy-
jsonb, jdbc-postgresql"
cd quartz-quickstart
```

It generates:

- the Maven structure
- a landing page accessible on <http://localhost:8080>
- example **Dockerfile** files for both **native** and **jvm** modes
- the application configuration file
- an **org.acme.quartz.TaskResource** resource
- an associated test

The Maven project also imports the Quarkus Quartz extension.

Creating the Task Entity

In the **org.acme.quartz** package, create the **Task** class, with the following content:

```

package org.acme.quartz;

import javax.persistence.Entity;
import java.time.Instant;
import javax.persistence.Table;

import io.quarkus.hibernate.orm.panache.PanacheEntity;

@Entity
@Table(name="TASKS")
public class Task extends PanacheEntity { ❶
    public Instant createdAt;

    public Task() {
        createdAt = Instant.now();
    }

    public Task(Instant time) {
        this.createdAt = time;
    }
}

```

1. Declare the entity using [Panache](#)

Creating a scheduled job

In the `org.acme.quartz` package, create the `TaskBean` class, with the following content:

```

package org.acme.quartz;

import javax.enterprise.context.ApplicationScoped;

import javax.transaction.Transactional;

import io.quarkus.scheduler.Scheduled;

@ApplicationScoped ❶
public class TaskBean {

    @Transactional
    @Scheduled(every = "10s") ❷
    void schedule() {
        Task task = new Task(); ❸
        task.persist(); ❹
    }
}

```

1. Declare the bean in the *application* scope
2. Use the `@Scheduled` annotation to instruct Quarkus to run this method every 10 seconds.
3. Create a new `Task` with the current start time.
4. Persist the task in database using `Panache`.

Updating the application configuration file

Edit the `application.properties` file and add the below configuration:

```
# Quartz configuration
quarkus.quartz.clustered=true ①
quarkus.quartz.store-type=db ②

# Datasource configuration.
quarkus.datasource.url=jdbc:postgresql://localhost/quarkus_test
quarkus.datasource.driver=org.postgresql.Driver
quarkus.datasource.username=quarkus_test
quarkus.datasource.password=quarkus_test

# Hibernate configuration
quarkus.hibernate-orm.database.generation=none
quarkus.hibernate-orm.log.sql=true
quarkus.hibernate-orm.sql-load-script=no-file

# flyway configuration
quarkus.flyway.connect-retries=10
quarkus.flyway.table=flyway_quarkus_history
quarkus.flyway.migrate-at-start=true
quarkus.flyway.baseline-on-migrate=true
quarkus.flyway.baseline-version=1.0
quarkus.flyway.baseline-description=Quartz
```

1. Indicate that the scheduler will be run in clustered mode
2. Use the database store to persist job related information so that they can be shared between nodes

Updating the resource and the test

Edit the `TaskResource` class, and update the content to:

```
package org.acme.quartz;

import java.util.List;

import javax.ws.rs.GET;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import javax.ws.rs.core.MediaType;

@Path("/tasks")
@Produces(MediaType.APPLICATION_JSON)
public class TaskResource {

    @GET
    public List<Task> listAll() {
        return Task.listAll(); ①
    }
}
```

1. Retrieve the list of created tasks from the database

We also need to update the tests. Edit the `TaskResourceTest` class to match:

```

package org.acme.quartz;

import io.quarkus.test.junit.QuarkusTest;

import static org.hamcrest.Matchers.greaterThanOrEqualTo;

import org.junit.jupiter.api.Test;

import static io.restassured.RestAssured.given;
import static org.hamcrest.CoreMatchers.is;

@QuarkusTest
public class TaskResourceTest {

    @Test
    public void tasks() throws InterruptedException {
        Thread.sleep(1000); // wait at least a second to have the
first task created
        given()
            .when().get("/tasks")
            .then()
            .statusCode(200)
            .body("size()", is(greaterThanOrEqualTo(1))); ①
    }
}

```

1. Ensure that we have a 200 response and at least one task created

Creating Quartz Tables

Add a SQL migration file named `src/main/resources/db/migration/V2.0.0QuarkusQuartzTasks.sql` with the content copied from file with the content from <https://github.com/quarkusio/quarkus-quickstarts/blob/master/quartz-quickstart/src/main/resources/db/migration/V2.0.0QuarkusQuartzTasks.sql> [V2.0.0__QuarkusQuartzTasks.sql].

Configuring the load balancer

In the root directory, create a `nginx.conf` file with the following content:

```
user  nginx;

events {
    worker_connections  1000;
}

http {
    server {
        listen 8080;
        location / {
            proxy_pass http://tasks:8080; ①
        }
    }
}
```

1. Route all traffic to our tasks application

Setting Application Deployment

In the root directory, create a `docker-compose.yml` file with the following content:

```

version: '3'

services:
  tasks: ①
    image: quarkus-quickstarts/quartz:1.0
    build:
      context: ./
      dockerfile: src/main/docker/Dockerfile.${QUARKUS_MODE:-jvm}
    environment:
      QUARKUS_DATASOURCE_URL:
jdbc:postgresql://postgres/quarkus_test
    networks:
      - tasks-network
    depends_on:
      - postgres

  nginx: ②
    image: nginx:1.17.6
    volumes:
      - ./nginx.conf:/etc/nginx/nginx.conf:ro
    depends_on:
      - tasks
    ports:
      - 8080:8080
    networks:
      - tasks-network

  postgres: ③
    image: postgres:11.3
    container_name: quarkus_test
    environment:
      - POSTGRES_USER=quarkus_test
      - POSTGRES_PASSWORD=quarkus_test
      - POSTGRES_DB=quarkus_test
    ports:
      - 5432:5432
    networks:
      - tasks-network

networks:
  tasks-network:
    driver: bridge

```

1. Define the tasks service
2. Define the nginx load balancer to route incoming traffic to an appropriate node
3. Define the configuration to run the database

Running the database

In a separate terminal, run the below command:

```
docker-compose up postgres ①
```

1. Start the database instance using the configuration options supplied in the `docker-compose.yml` file

Run the application in Dev Mode

Run the application with: `./mvnw quarkus:dev`. After a few seconds, open another terminal and run `curl localhost:8080/tasks` to verify that we have at least one task created.

As usual, the application can be packaged using `./mvnw clean package` and executed using the `-runner.jar` file. You can also generate the native executable with `./mvnw clean package -Pnative`.

Packaging the application and run several instances

The application can be packaged using `./mvnw clean package`. Once the build is successful, run the below command:

```
docker-compose up --scale tasks=2 --scale nginx=1 ①
```

1. Start two instances of the application and a load balancer





After a few seconds, in another terminal, run `curl localhost:8080/tasks` to verify that tasks were only created at different instants and in an interval of 10 seconds.

You can also generate the native executable with `./mvnw clean package -Pnative`.

Quartz Configuration Reference

🔒 Configuration property fixed at build time - ⚙️ Configuration property overridable at runtime

Configuration property	Type	Default
🔒 <code>quarkus.quartz.clustered</code>		
Enable cluster mode or not. If enabled make sure to set the appropriate cluster properties.	boolean	<code>false</code>

 <code>quarkus.quartz.store-type</code> <p>The type of store to use. When using the <code>db</code> store type configuration value make sure that you have the datasource configured. See Configuring your datasource for more information. To create Quartz tables, you can perform a schema migration via the Flyway extension using a SQL script matching your database picked from Quartz repository.</p>	<p>ram, db</p>	<p>ram</p>
 <code>quarkus.quartz.datasource</code> <p>The name of the datasource to use. Optionally needed when using the <code>db</code> store type. If not specified, defaults to using the default datasource.</p>	<p>string</p>	
 <code>quarkus.quartz.thread-count</code> <p>The size of scheduler thread pool. This will initialize the number of worker threads in the pool.</p>	<p>int</p>	<p>25</p>
 <code>quarkus.quartz.thread-priority</code> <p>Thread priority of worker threads in the pool.</p>	<p>int</p>	<p>5</p>