

Docker instructions

What is Docker?

Docker is a software, which enables the containerization of applications at an OS-level. This way containerized applications can be run on any device running Docker.

For further information see: <https://www.docker.com/>

Get started

First, install Docker. An easy way to install Docker is to **install using the convenience script** provided for various Linux distributions, e.g. Ubuntu: <https://docs.docker.com/engine/install/ubuntu/>

After installing Docker, you can test its functionality by executing the following command in the terminal:

```
sudo docker run docker/whalesay cowsay Hello World
```

This is the Docker counterpart to a “Hello world” program. It will pull the corresponding Docker image from Dockerhub (the central image repository) and then run it with the specified string.

How to use the provided Dockerfile

Dockerfiles basically contain a set of instructions, which Docker uses to build an image. The provided Dockerfile uses Ubuntu:20.04 as a base image and installs the Java JDK 11, Python including psycopg2 and PostgreSQL, when building an image from it.

To build an image from the Dockerfile simply execute the following command in the same directory:

```
sudo docker build -t <my-postgres-image> .
```

After this, you can run a container from the image:

```
sudo docker run -d --name <my-postgres-container> <my-postgres-image>
```

On build, Docker will copy all files from the current directory and copy them into the working directory of the container. Additionally, it will create a database ‘mydb’ by user ‘docker’ with password ‘admin’.

To enter the shell of running container, execute:

```
sudo docker exec -it <my-postgres-container> /bin/bash
```

Note: If a CreateSchema.sql was in the same directory as the Dockerfile, it will have been copied into the working directory of the Docker container. By executing the following command, you can run a SQL script on any database, in this case ‘mydb’:

```
psql -d <mydb> -a -f <CreateSchema.sql>
```

To execute queries on a database, simply execute the following commands:

```
psql
```

```
\c <mydb>
```

Lastly, you can execute Java or Python programs just like on any other Linux machine, e.g.:

```
python3 IngestData.py [..]
```

with the corresponding command line arguments.

Additional docker commands:

sudo docker ps

show all running containers

sudo docker kill <my-postgres-container>

kill a running container

sudo docker stop <my-postgres-container>

stop a running container

sudo docker start <my-postgres-container>

restart a container

sudo docker rm <my-postgres-container>

remove a container

sudo docker rmi <my-postgres-image>

remove an image