Docker Image

`docker images`	See all available images
`docker build -t <na me>`	Create an image with a pretty name (you must define the `Do ckerfile` in the folder)
`docker tag <na me> userna me/ rep osi tor y:tag`	This tags an image ready to be sent to a repository
`docker push userna me/ rep osi tor y:tag`	Push the image to the remote repository
`docker search <ke ywo rd>`	Search for public repos- itories

Docker Images are the base for `containers` and are similar to `.is o` files. They can be for example the image of your app and contain everything needed to run the application.

These images can be local or in repositories (and marked with an tag)

To create images, you must create a `Dockerfile` with some docker commands to specify how that image will be created, for example to setup the environment and a BaseImage.

Services

Services (cont)

To run it you must first start: `docker swarm init`
Then run it giving a name:
`docker stack deploy -c docker -co mpo se.yml <ap p_n ame>`
To see the details of containers running in your service, run:
`docker stack ps <ap p_n ame>`
Now each time you request your app (via browser, for example), the load-balancer will help you distributing the requests to each replica.
To put it down, `docker stack rm <ap p_n ame>`
This puts down the app, but not the "one-node" swarm we created. Use
`docker swarm leave --force`

Docker Swarm is available only for version "3"

Docker Container

`docker run <im age>`	Run the image, starting a Container
`-d`	Run in detached mode (in background)
`-p 4000:80`	Maps the port `80` of the image to the host port `4000`
`--rm`	Removes the container when exited
`docker ps`	List the running containers (you can check container id)
`docker ps -l`	List all the containers (you can check container id)
`docker stop <co nta ine r_ i d>`	Stop the container

When you run an image with you are starting a Container, so container is the runtime instance of an `image`, and consists of the image, an execution environment and a standart set of instructions.

Swarm

`docker swarm init`	Initialize a swarm and become `swarm manager`
`docker swarm join`	Join a swarm as `worker`
`docker swarm leave --force`	Leaves the current swarm

With Docker you can increase resource and capacities by creating a `swarm`, which are simply several machines (virtual or physical) running a Docker and joined to a cluster.

Swarms have the `swarm manager`, which can issue docker commands normally, and the `workers` which are only there to provide capacity.

Different pieces of the app are called "services" For example, a service for storing application data in a database, a service for the front-end, etc.

Services are just "containers in production." A service only runs one image, but it manages for example what ports it should use and how many replicas of the container should run.

To define a service, you'll need an `docker -co mpo se.yml` file. For example:

```
version: " 3"
services:
  web:
    image: amiche let ti/ get -st art ed: part1
    deploy:
      replicas: 5
      resources:
        limits:
          cpus: " 0.1 "
          memory: 50M
      restar t_p olicy:
        condition: on-failure
    ports:
      - " 80: 80"
    networks:
      - webnet
networks:
  webnet:
```

Here you define the image to be loaded, how many replicas, the resource limits and the restart conditions.