

# An overview of some container-based technologies

Containerization workshop

Infrastructures for Cloud Computing and Big Data M

Dmitrij David Padalino Montenero

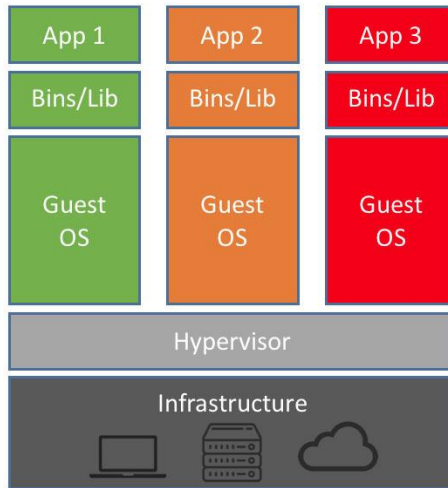
Academic Year 2018/2019

- Containerization and Orchestration
- Docker
- Kubernetes
- OpenShift
- IBM Cloud Kubernetes-as-a-Service

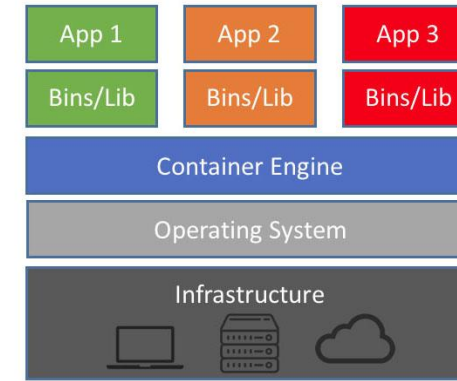
**Containerization** refers to an operating system feature in which the kernel allows the existence of multiple isolated user-space instances.

Such instances, called **containers**, partitions or virtual environments may look like real computers from the point of view of programs running in them.

# Virtual Machines vs Containers - 1

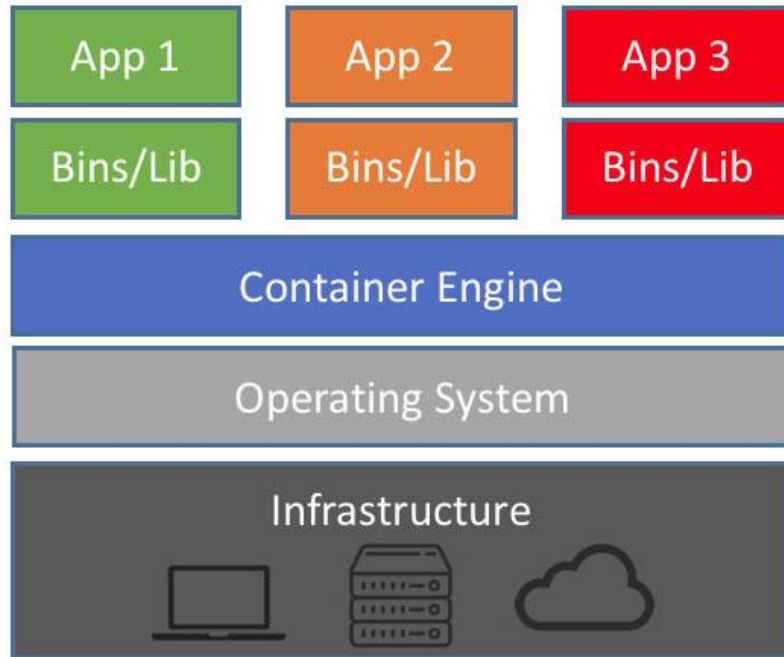


- Hypervisor
- Many guest OSs in the same physical server
- Visibility of all hardware resources
- Well separated run-time environments
- Images are Gigabyte in size
- Hardware resources optimization
- Cost savings



- Container engine
- Host OS kernel sharing
- Many containers in the same physical server
- Visibility only of the resources assigned to containers
- Relatively well separated run-time environments
- Images are Megabytes in size
- Further hardware resources optimization

# Virtual Machines vs Containers - 2



Novel way to think about the architecture of a distributed application, so called **micro-services architecture**.

Each container is usually single-purpose, single-process, which aligns nicely with micro-services architectures.

The new approach is to package the different services that constitute an application into separate containers, and to deploy those containers across a cluster of physical or virtual machines.

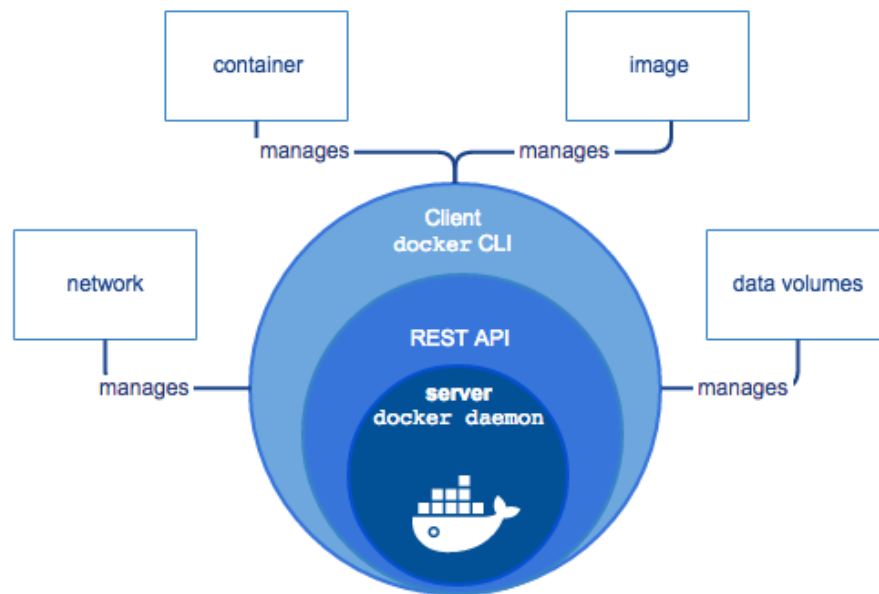
Containers are

- Flexible
- Lightweight
- Interchangeable
- Portable
- Scalable

- Need to manage micro-services applications.
- **Container orchestration** allows users to control when containers start and stop, group them into clusters, and coordinate all of the processes that compose an application. Container orchestration tools allow users to guide container deployment and automate updates, health monitoring, and failover procedures.

Docker is a platform for developers and sysadmins to **develop**, **deploy**, and **run** applications with containers.

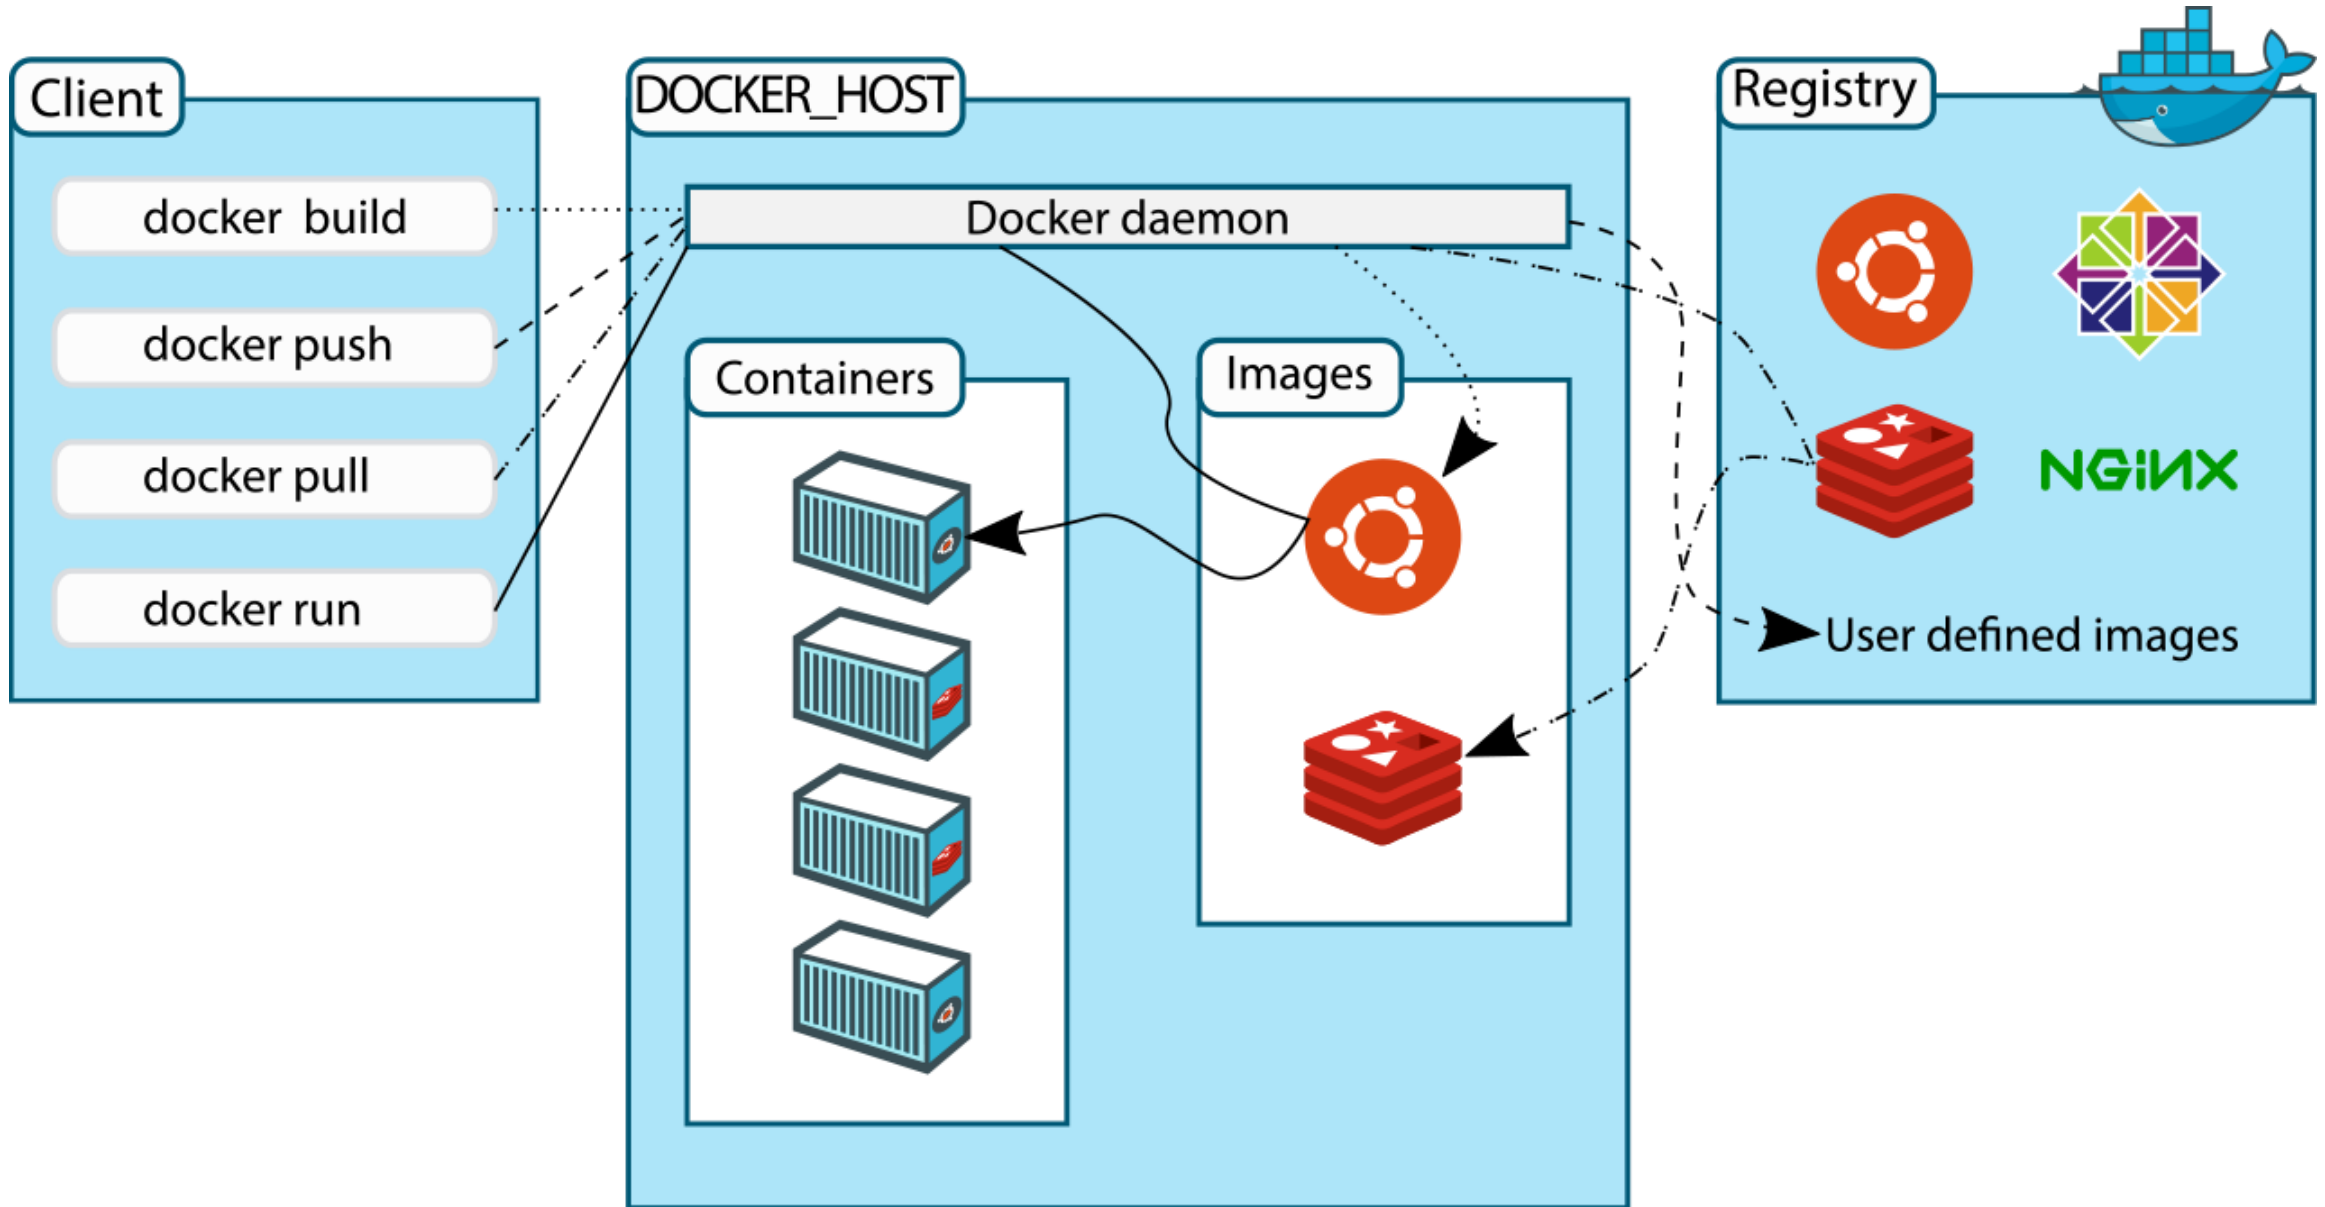
It enables you to separate your applications from your infrastructure so you can deliver software quickly.



## Client-server architecture

- Server: docker daemon
- API REST
- Client: docker CLI

# Architecture





## Starting from a 3-tier web application (Spring and MySQL)

- Application containerization
- Creation and configuration of a cluster (swarm) with Docker Machine.
- Deployment.
- Horizontal scaling.

CPU	Intel(R) Core(TM) i7-3610QM CPU @ 2.30Ghz
RAM	8GB
SSD	Samsung 850 Evo 500GB
Interfaccia di rete wireless	Qualcomm Atheros AR5BWB22
Sistema operativo	Ubuntu 18.04.1 LTS Bionic Beaver
<b>Versione Oracle VM VirtualBox</b>	5.2.20r125813
<b>Versione Docker</b>	18.06.1-ce
<b>Versione Docker Machine</b>	0.14.0
Versione Google Chrome	69.0.3497.100 a 64 bit

Source code link: <https://github.com/davidMonnuar/projectActivity>

## 3-tier web application for books collection management

### Client Tier



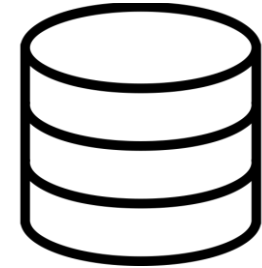
- HTML5
- Javascript
- CSS

### Application Tier

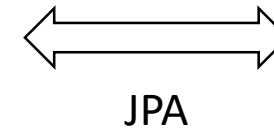


- Java
- Spring Boot
- Tomcat

### Database Tier



- Rel. DB
- MySQL



**Spring Boot** is an approach to develop spring application with minimal or zero configuration.

## **Tomcat used as an Embedded Server**

Think about what you would need to be able to deploy your application (typically) on a virtual machine.

Step 1: Install Java

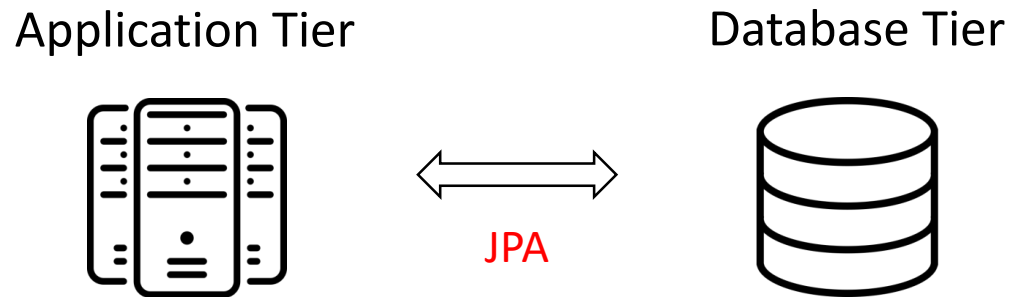
Step 2: Install the Web/Application Server (Tomcat/Websphere/Weblogic etc)

Step 3: Deploy the application war

What if we want to simplify this? This idea is the genesis for Embedded Servers.

For example, for a Spring Boot Application, you can generate an application jar which contains Embedded Tomcat. **You can run a web application as a normal Java application!**

Thanks to **Spring deep integration with JPA** we don't need to spend too much time in configuring the persistence layer since most part of the setup is done automatically (i.e. db creation) by exploiting Java annotations.



The only thing we have to write down is a configuration file called **application.yml** in which we specify

- the URL to access the database
- username
- password
- plus other parameters

## Entities

- Author.java
- Book.java

## Repositories

- AuthorRepository.java
- BookRepository.java

## Controller

- MainController.java

## View

- CustomBookSerializer.java

## Resources

- application.yml

```
package it.unibo.libraryDemo.entities;

import com.fasterxml.jackson.databind.annotation.JsonSerialize;
import it.unibo.libraryDemo.view.CustomBooksSerializer;

import javax.persistence.*;
import java.io.Serializable;
import java.util.ArrayList;
import java.util.List;

@Entity
@Table(name="authors")
public class Author implements Serializable {
    @Id
    @Column(name="author_id")
    @GeneratedValue(strategy=GenerationType.AUTO)
    private Integer id;

    private String name;

    @OneToMany(
        mappedBy = "author",
        cascade = CascadeType.ALL,
        orphanRemoval = true
    )
    @JsonSerialize(using = CustomBooksSerializer.class)
    private List<Book> books = new ArrayList<>();

    @PreRemove
    public void preRemove() {
        setBooks(null);
    }
}
```

```
public String getName() {
    return name;
}

public void setName(String name) {
    this.name = name;
}

public Integer getId() {
    return id;
}

public void setId(Integer id) {
    this.id = id;
}

public List<Book> getBooks() {
    return books;
}

public void setBooks(List<Book> books) {
    this.books = books;
}

public void addBook(Book book) {
    books.add(book);
    book.setAuthor(this);
}

public void removeBook(Book book) {
    books.remove(book);
    book.setAuthor(null);
}
}
```

```
package it.unibo.libraryDemo.repositories;

import it.unibo.libraryDemo.entities.Author;
import org.springframework.data.repository.CrudRepository;

public interface AuthorRepository extends CrudRepository<Author, Integer> {

}
```

**CrudRepository** is an interface and extends Spring data **Repository** interface.

CrudRepository provides generic CRUD operation on a repository for a specific type.

It has generic methods for CRUD operation.

To use CrudRepository we have to create our interface and extend CrudRepository. We need not to implement our interface, **its implementation will be created automatically at runtime**. Some methods we will use

- findById
- findAll
- deleteById
- save

In this class are defined and implemented the REST APIs that will be used by the HTML5 Client.

The APIs are

- library/**getHostname**
- library/**resetApplication**
- library/**addBook**
- library/**getAllBooks**
- library/**deleteBook**
- library/**addAuthor**
- library/**getAllAuthors**
- library/**deleteAuthor**

Client Tier



Application Tier



```
@RestController
@RequestMapping(path="/library")
public class MainController {
    @Autowired
    private AuthorRepository authorRepository;
    @Autowired
    private BookRepository bookRepository;

    @GetMapping(path="/getHostname")
    public String getHostname () {
        String hostname = "unknown";
        try {
            hostname = InetAddress.getLocalHost().getHostName();
        } catch (UnknownHostException e) {
            e.printStackTrace();
        }
        return hostname;
    }

    @GetMapping(path="/resetApplication")
    public String resetApplication () {
        bookRepository.deleteAll();
        authorRepository.deleteAll();
        return "Done";
    }

    @PutMapping(path="/addBook")
    public String addNewBook (@RequestParam String title,
        @RequestParam String isbn, @RequestParam String author) {
        Book b = new Book();
        b.setTitle(title);
        b.setIsbn(isbn);
        Author a = authorRepository.
            findById(Integer.parseInt(author)).orElse(null);
        if(a == null) {
            return "Not saved. The author is not in the database";
        }
    }
```

```
        a.addBook(b);
        bookRepository.save(b);
        authorRepository.save(a);
        return "Book saved";
    }

    @GetMapping(path="/getAllBooks")
    public Iterable<Book> getAllBooks () {
        return bookRepository.findAll();
    }

    @DeleteMapping(path="/deleteBook")
    public String deleteBook (@RequestParam Integer id){
        bookRepository.deleteById(id);
        return "Book deleted";
    }

    @PutMapping(path="/addAuthor")
    public String addNewAuthor (@RequestParam String name) {
        Author a = new Author();
        a.setName(name);
        authorRepository.save(a);

        return "Author saved";
    }

    @GetMapping(path="/getAllAuthors")
    public Iterable<Author> getAllAuthors () {
        return authorRepository.findAll();
    }

    @DeleteMapping(path="/deleteAuthor")
    public String deleteAuthor (@RequestParam Integer id){
        authorRepository.deleteById(id);
        return "Author deleted";
    }
}
```



We use global variables in order to avoid to hard-code the database information.

This let us to re-use this web app with any database we want.

```
spring:
  profiles: container
  datasource:
    driverClassName: ${DATABASE_DRIVER}
    url: jdbc:mysql://${DATABASE_HOST}:${DATABASE_PORT}/${DATABASE_NAME}?useSSL=false&allowPublicKeyRetrieval=true
    username: ${DATABASE_USER}
    password: ${DATABASE_PASSWORD}
  tomcat:
    test-while-idle: true
    time-between-eviction-runs-millis: 60000
    validation-query: SELECT 1
  jpa:
    hibernate.ddl-auto: create-drop
    properties.hibernate.dialect: org.hibernate.dialect.MySQL5Dialect
```

## App organization

- One **container** for Spring
- One **container** for MySQL + **Volume**
- One **overlay network** for connecting the two containers

The first step is to define an image for each container by writing a text document called **Dockerfile**.

Thanks to a simple syntax a Dockerfile let us to define all the steps necessary to create an image and to execute it.

## Spring container Dockerfile

```
1 FROM openjdk:8-jre-alpine
2 MAINTAINER Dmitrij David Padalino Montenero
3
4 VOLUME /tmp
5 EXPOSE 8080
6
7 ARG JAR_FILE
8 ADD target/${JAR_FILE} app.jar
9 ADD wrapper.sh wrapper.sh
10
11 RUN apk add --update bash && rm -rf /var/cache/apk/*
12 RUN bash -c 'chmod +x /wrapper.sh'
13 RUN bash -c 'touch /app.jar'
14
15 ENTRYPOINT ["/bin/bash", "/wrapper.sh"]
```

## Script entrypoint

```
1 #!/bin/bash
2 while ! exec 6<>/dev/tcp/${DATABASE_HOST}/${DATABASE_PORT}; do
3 echo "Trying to connect to MySQL at ${DATABASE_PORT}..."
4 sleep 10
5 done
6
7 java -Djava.security.egd=file:/dev/./urandom
-Dspring.profiles.active=container -jar /app.jar
```

## Image creation and upload in the DockerHub

```
$ docker build -t librarydemo
$ docker login
$ docker tag librarydemo davidmonnuar/springbootlibrarydemo:1.0
$ docker push davidmonnuar/springbootlibrarydemo:1.0
```

Image size only 116 MB

## MySQL container Dockerfile

No need to create a new Docker image since the official MySQL image available in the DockerHub is enough for our purposes.

A **swarm** is a group of machines, physical or virtual, that are running Docker and joined into a cluster.

The cluster we will use is composed by two nodes, a master and one worker.

## Virtual machines creation

```
$ docker-machine create --driver virtualbox myvm1
```

```
$ docker-machine create --driver virtualbox myvm2
```

```
$ docker-machine ls
```

NAME	ACTIVE	DRIVER	STATE	URL	SWARM	DOCKER	ERRORS
myvm1	-	virtualbox	Running	tcp://192.168.99.100:2376		v18.06.1-ce	
myvm2	-	virtualbox	Running	tcp://192.168.99.101:2376		v18.06.1-ce	

## Cluster initialization

```
$ docker-machine ssh myvm1 "docker swarm init --advertise-addr 192.168.99.100:2377"
```

```
Swarm initialized: current node 0unmutpbeeeytxpquhiylb6ok is now a manager.
```

```
To add a worker to this swarm, run the following command:
```

```
docker swarm join \
```

```
--token SWMTKN-1-4y8bpxyrbno89dapkkwyjdw3268qfzp128tpf5hecjdti5hb1k-00oatop19dzw3jejbieta9vp 192.168.99.100:2377
```

```
$ docker-machine ssh myvm2 "docker swarm join --token SWMTKN-1-4y8bpxyrbno89dapkkwyjdw3268qfzp128tpf5hecjdti5hb1k-00oatop19dzw3jejbieta9vp 192.168.99.100:2377"
```

```
$ docker-machine ssh myvm1 "docker node ls"
```

ID	HOSTNAME	STATUS	AVAILABILITY	MANAGER STATUS	ENGINE VERSION
0unmutpbeeeytxpquhiylb6ok *	myvm1	Ready	Active	Leader	18.06.1-ce
ktvyi0nypz31l645iej8torkx	myvm2	Ready	Active		18.06.1-ce

In a distributed application, different pieces of the app are called “services”.

**Services** in Docker are really just “containers in production.” A service only runs one image, but it codifies the way that image runs - what ports it should use, how many replicas of the container should run so the service has the capacity it needs, and so on. Scaling a service changes the number of container instances running that piece of software, assigning more computing resources to the service in the process.

In order to connect services Docker uses the concept of **Stack**. A stack is a group of interrelated services that share dependencies, and can be orchestrated and scaled together.

In the **docker-compose.yml** file are defined

- A service “**mysql**db”, port 3306
- A service “**spring**”, port 80
- A service “**visualizer**”, port 8080
- A volume “**my-data**volume” used by “mysql”db”
- A network “**web**net” with the default settings (which is a load-balanced overlay network)

## Application deployment

```
$ docker-machine ssh myvm1 "docker stack deploy -c docker-compose.yml librarywebapp"
```

```
1  version: "3"
2  services:
3    mysqldb:
4      image: mysql:latest
5      volumes:
6        - my-datavolume:/var/lib/mysql
7      deploy:
8        placement:
9          constraints: [node.role == manager]
10     environment:
11       - MYSQL_ROOT_PASSWORD=password
12       - MYSQL_DATABASE=db_example
13       - MYSQL_USER=springuser
14       - MYSQL_PASSWORD=mysql2019
15     networks:
16       - webnet
17   spring:
18     image: davidmonnuar/springbootlibrarydemo:1.0
19     depends_on:
20       - mysqldb
21     deploy:
22       replicas: 5
23       restart_policy:
24         condition: on-failure
25       resources:
26         limits:
27           cpus: "0.15"
28           memory: 500M
29     ports:
30       - "80:8080"
31
```

```
31  environment:
32    - DATABASE_DRIVER=com.mysql.jdbc.Driver
33    - DATABASE_HOST=mysqldb
34    - DATABASE_PORT=3306
35    - DATABASE_NAME=db_example
36    - DATABASE_USER=springuser
37    - DATABASE_PASSWORD=mysql2019
38  networks:
39    - webnet
40  visualizer:
41    image: dockersamples/visualizer:stable
42    ports:
43      - "8080:8080"
44    volumes:
45      - "/var/run/docker.sock:/var/run/docker.sock"
46    deploy:
47      placement:
48        constraints: [node.role == manager]
49    networks:
50      - webnet
51  volumes:
52    my-datavolume:
53  networks:
54    webnet:
```

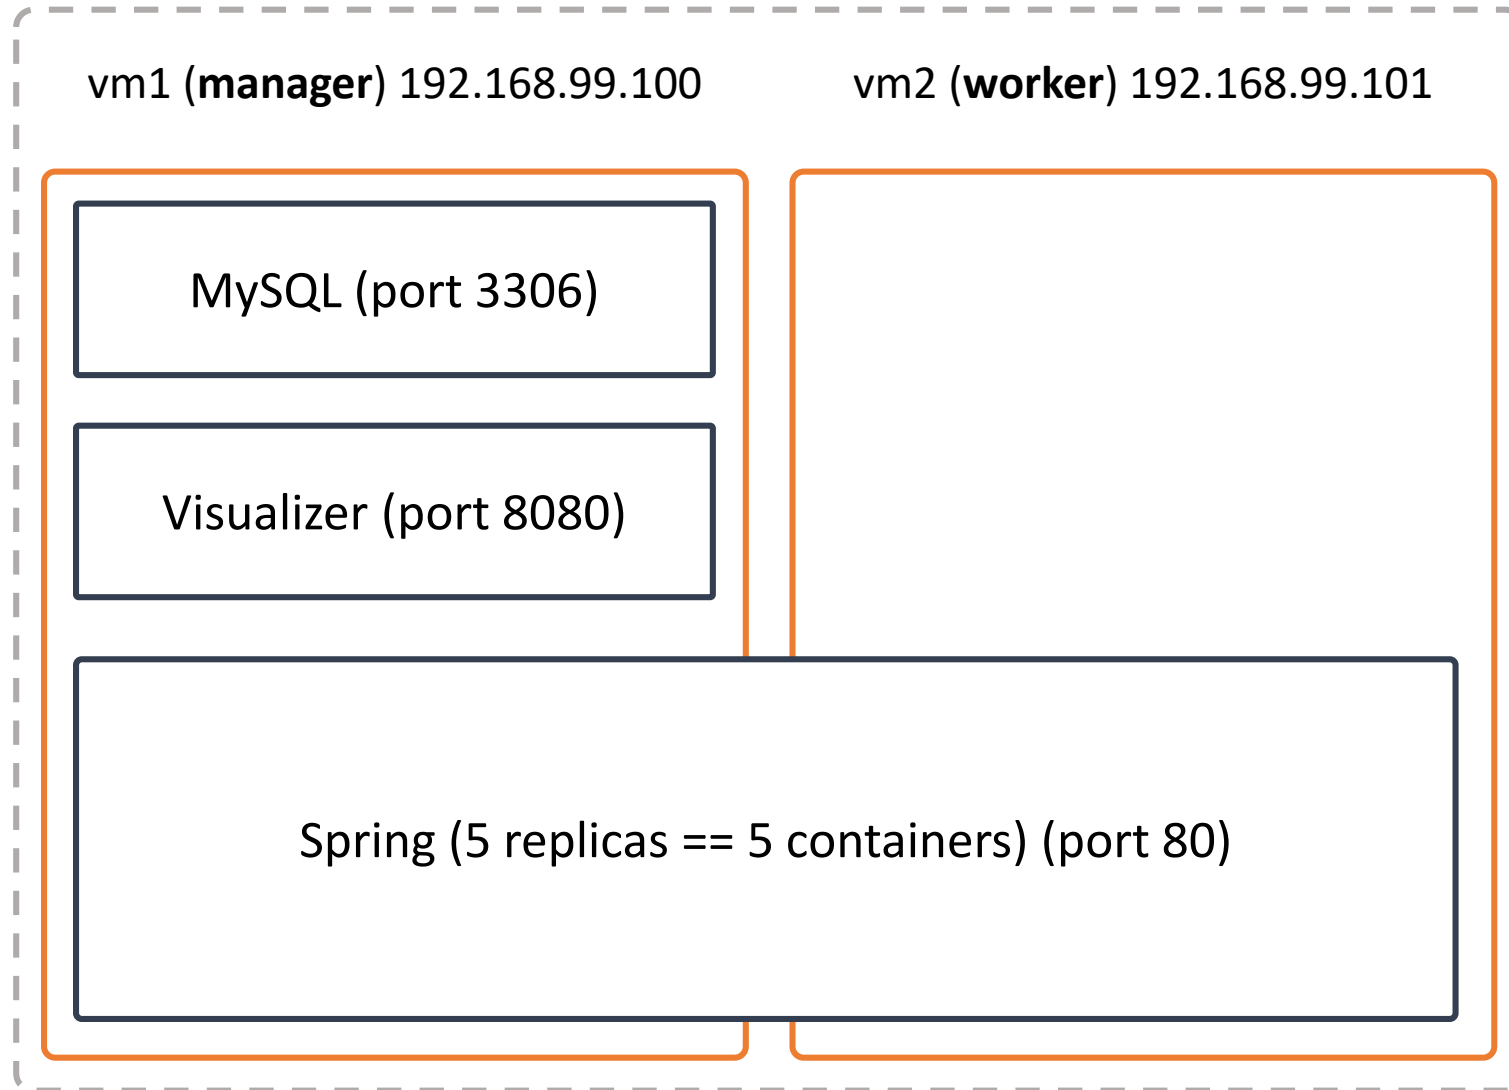
# Deployment - 3



-- Swarm

— Node

— Service



# Deployment - 4

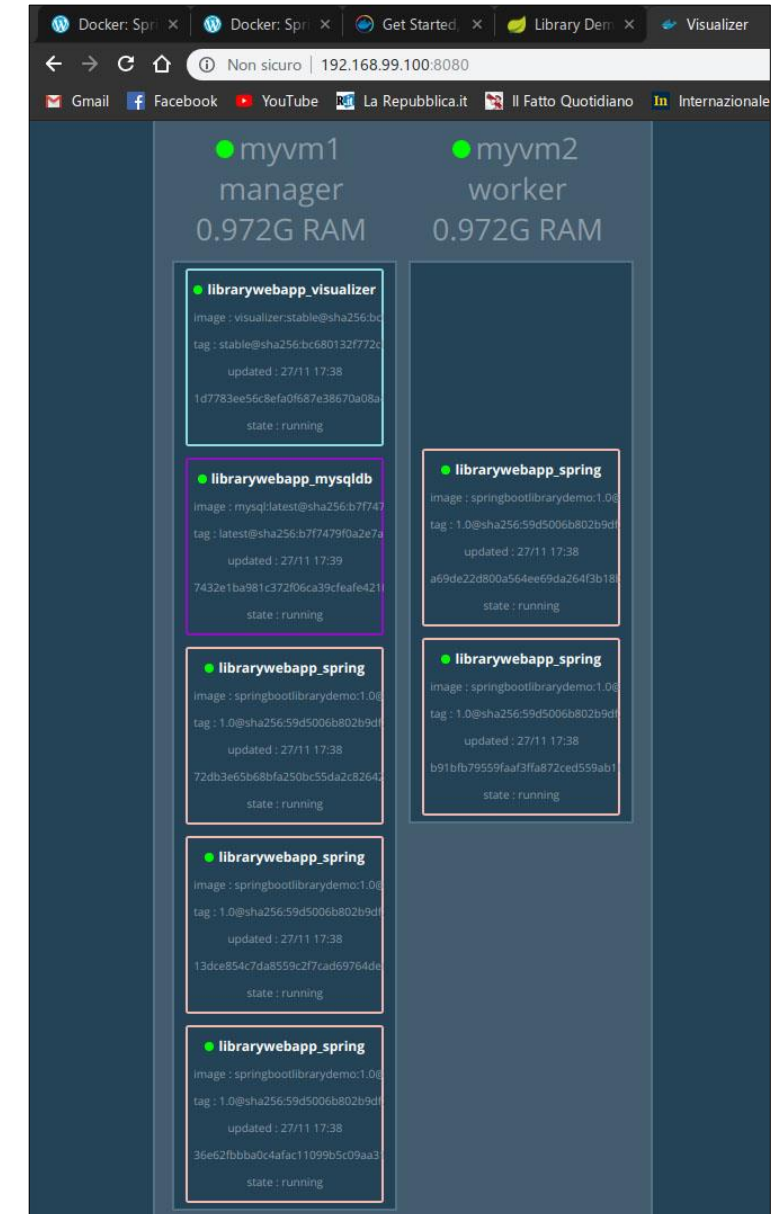
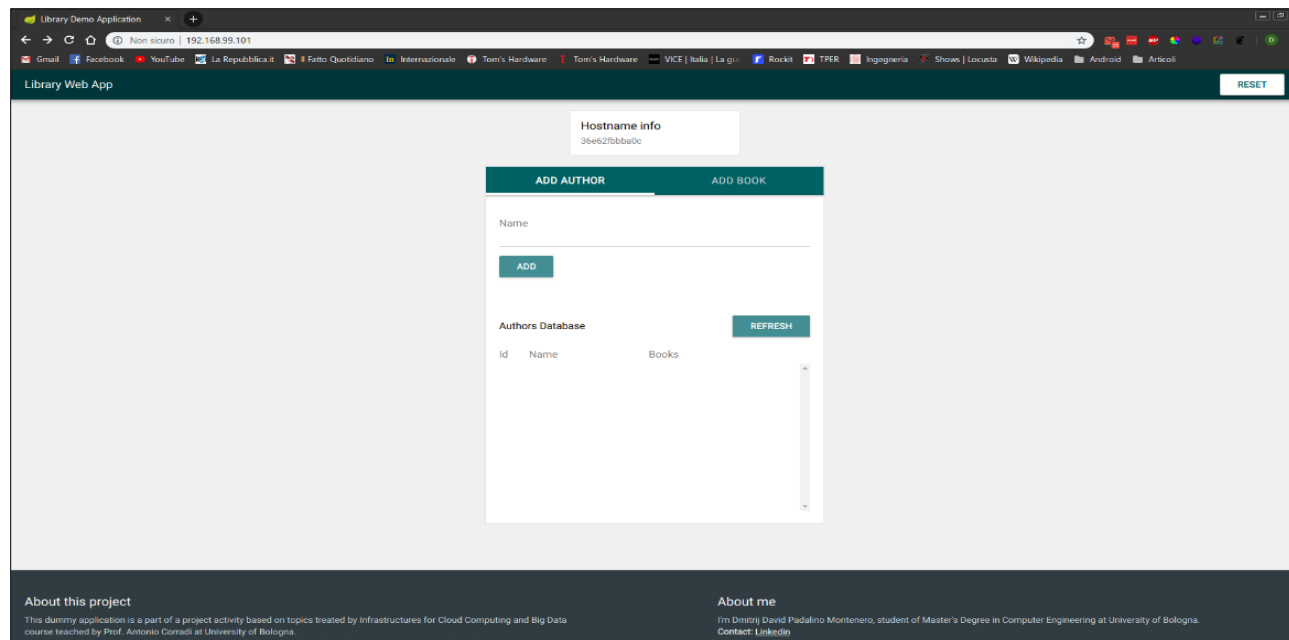


By inserting one of the two cluster node URLs in the browser it is possible to have **access to the application**

- <http://192.168.99.100>
- <http://192.168.99.101>

For the *Visualizer* service

- <http://192.168.99.100:8080>
- <http://192.168.99.101:8080>





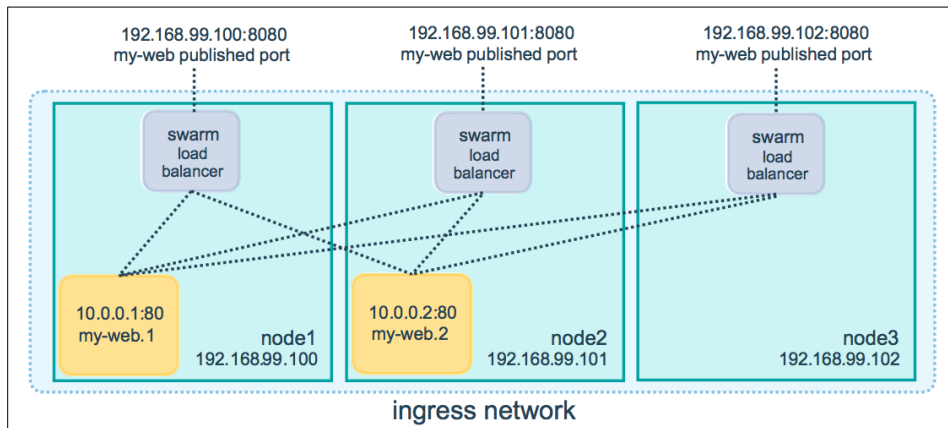
# Deployment - 5 and Scaling



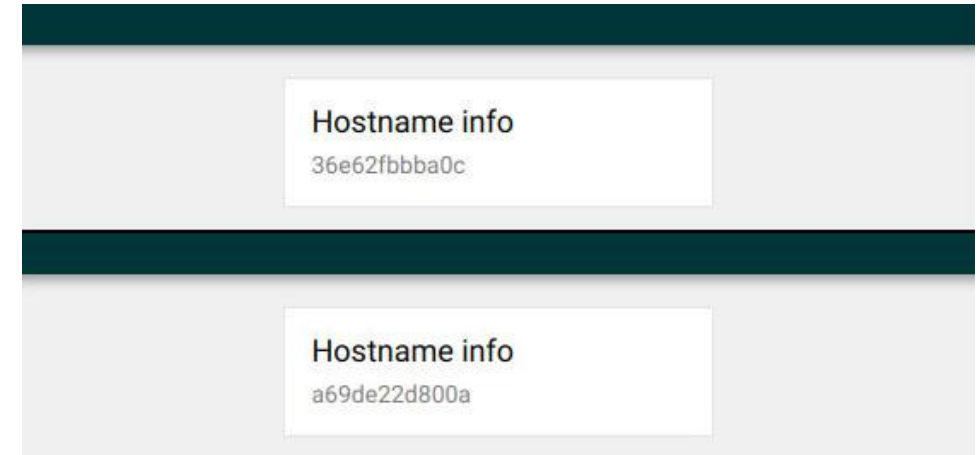
It is possible to access the app from the IP address of either myvm1 or myvm2

- `http://192.168.99.100`
- `http://192.168.99.101`

The reason both IP addresses work is that nodes in a swarm participate in an **ingress routing mesh**. This ensures that a service deployed at a certain port within your swarm always has that port reserved to itself, no matter what node is actually running the container.



The overlay network *webnet* is load-balanced by using a round-robin strategy.



## Spring service horizontal scaling

You can scale the app by changing the replicas value in `docker-compose.yml`, saving the change, and re-running in the swarm manager

```
$ docker-machine ssh myvm1 "docker stack deploy -c docker-compose.yml librarywebapp"
```

Docker is a powerful tool easy to use. The documentation available is very rich and it can be found [here](#).

You've learned how to

- **containerize** an application by writing a Dockerfile of just few lines.
- **deploy** the app in a cluster
- **scale** the app

However to launch the application and scale it is necessary to access the cluster via ssh and by using Docker Swarm **is not possible** to define more fine-grained policies in order to dynamically scale the application.

From this experience it is interesting to understand how an orchestrator different from Docker Swarm, like **Kubernetes**, is able to tackle this problem.

Kubernetes is an open-source container orchestration system designed by Google for automating application deployment, scaling and management inside a cluster.

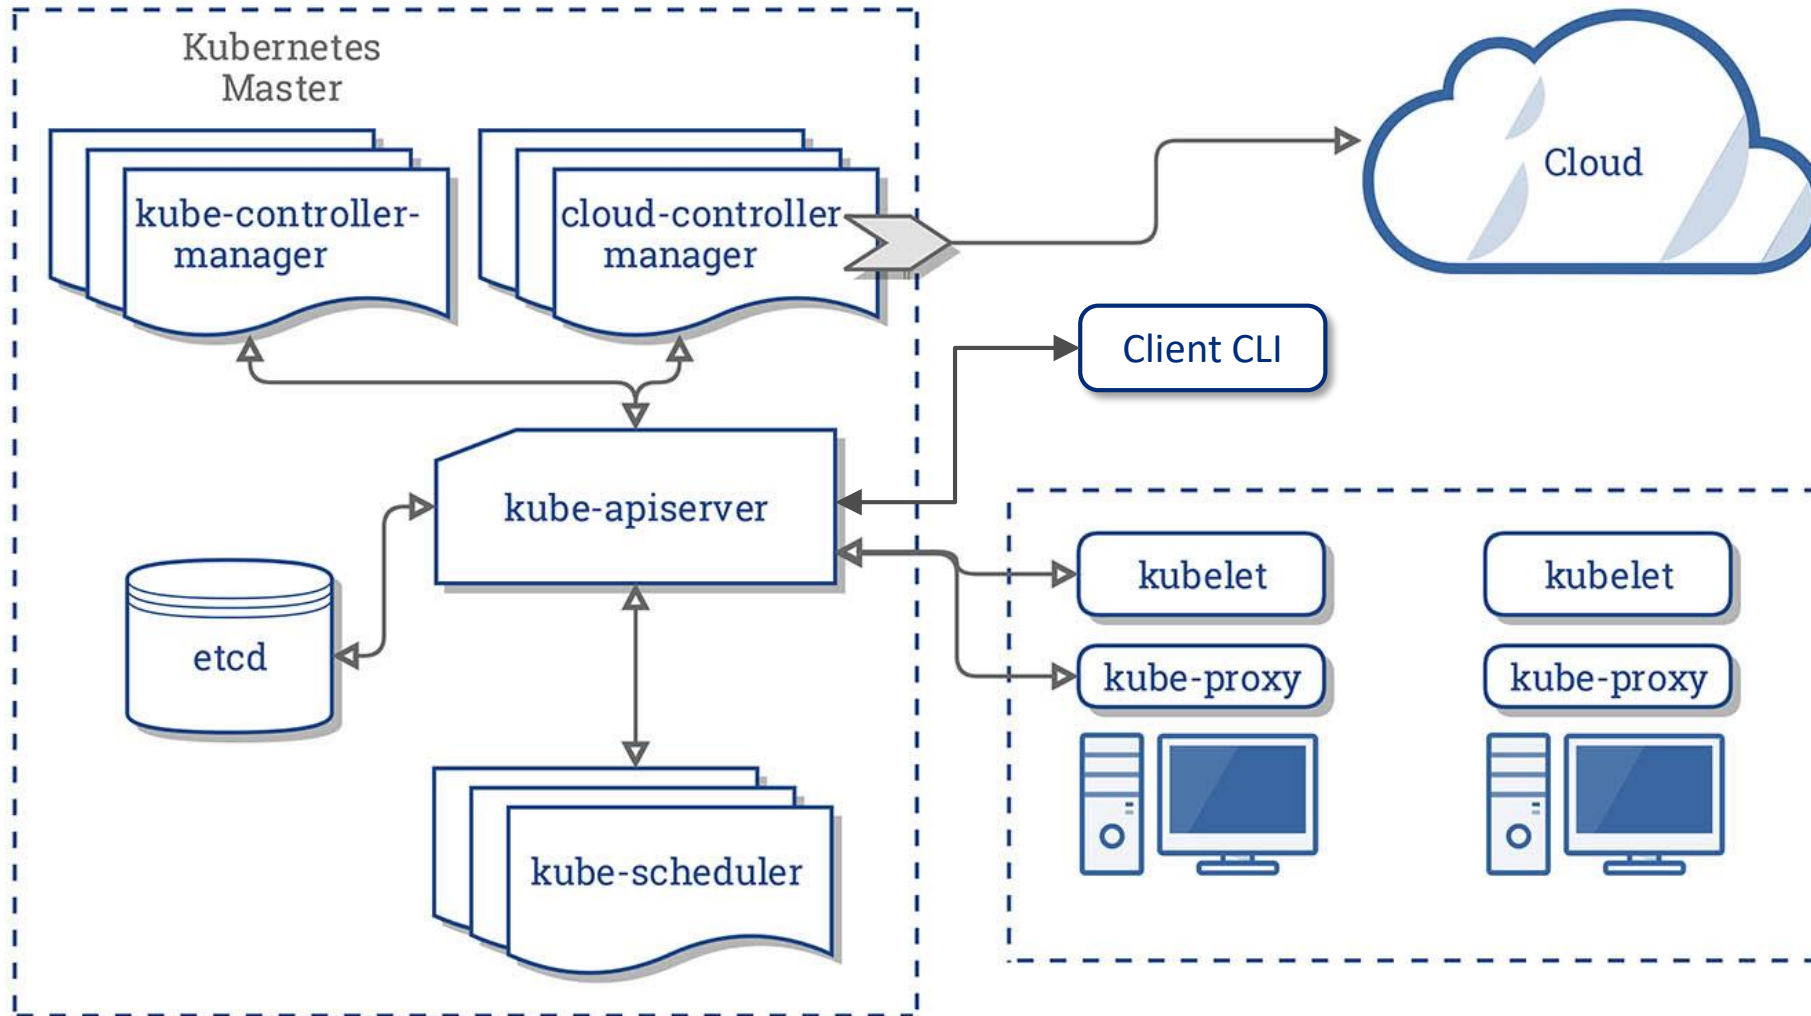
The **main goal** is to hide the complexity of managing a fleet of containers by providing to the user a set of REST APIs.

Kubernetes is **portable** in nature, meaning it can run on various public or private cloud platforms such as AWS, Azure, OpenStack or Apache Mesos.

- Multi-container application deployment.
- Scale apps in a automatic and dynamic way.
- Roll out updates to your app or its configuration by preserving service availability.
- It provides computational, network and storage resources to your app and a discovery service.
- Independency from the underlying infrastructure.

Kubernetes is more powerful than Docker Swarm, and requires more work to deploy.

However the work is intended to provide a big payoff in the long run in terms of a more manageable, resilient application infrastructure



## Master-slave architecture.

### Master node

- Etcd
- Kube-apiserver
- Kube-controller-manager
- Kube-scheduler

### Worker node

- Kubelet
- Kube-proxy

### Client

- CLI
- Dashboard

- **Pod**

represents a single instance of an application or running process in Kubernetes, and consists of one or more containers. Kubernetes starts, stops, and replicates all containers in a pod as a group.

- **Service**

because pods live and die as needed, we need a different abstraction for dealing with the application lifecycle. An application is supposed to be a persistent entity, even when the pods running the containers that comprise the application aren't themselves persistent. To that end, Kubernetes provides an abstraction called a service. A service describes how a given group of pods (or other Kubernetes objects) can be accessed via the network.

- **Volume**

similar to Docker Volumes, the main difference in Kubernetes is that the association between a volume and the container is done at the pod level. So the volume is attached to all containers running in that pod.

- **Namespace**

a virtual cluster (a single physical cluster can run multiple virtual ones) intended for environments with many users spread across multiple teams or projects, for isolation of concerns. Also, a namespace can be allocated a resource quota to avoid consuming more than its share of the physical cluster's overall resources.

Starting from the same web app (Spring and MySQL)

- Cluster creation and configuration with Minikube
- Deployment
- Static and dynamic horizontal scaling

CPU	Intel(R) Core(TM) i7-3610QM CPU @ 2.30Ghz
RAM	8GB
SSD	Samsung 850 Evo 500GB
Interfaccia di rete wireless	Qualcomm Atheros AR5BWB22
Sistema operativo	Ubuntu 18.04.1 LTS Bionic Beaver
<b>Versione Oracle VM VirtualBox</b>	5.2.20r125813
<b>Versione Minikube</b>	0.30.0
<b>Versione kubectl Client</b>	1.12.3
<b>Versione kubectl Server</b>	1.10.0
Versione Google Chrome	69.0.3497.100 a 64 bit

Minikube is a tool that makes it easy to run Kubernetes locally.  
It runs a single-node Kubernetes cluster inside a VM.

## Cluster initialization

```
$ minikube start
```

```
$ kubectl cluster-info
```

```
Kubernetes master is running at https://192.168.99.100:8443
```

```
CoreDNS is running at https://192.168.99.100:8443/api/v1/namespaces/kubedns/services/kube-dns:dns/proxy
```

```
$ kubectl get nodes
```

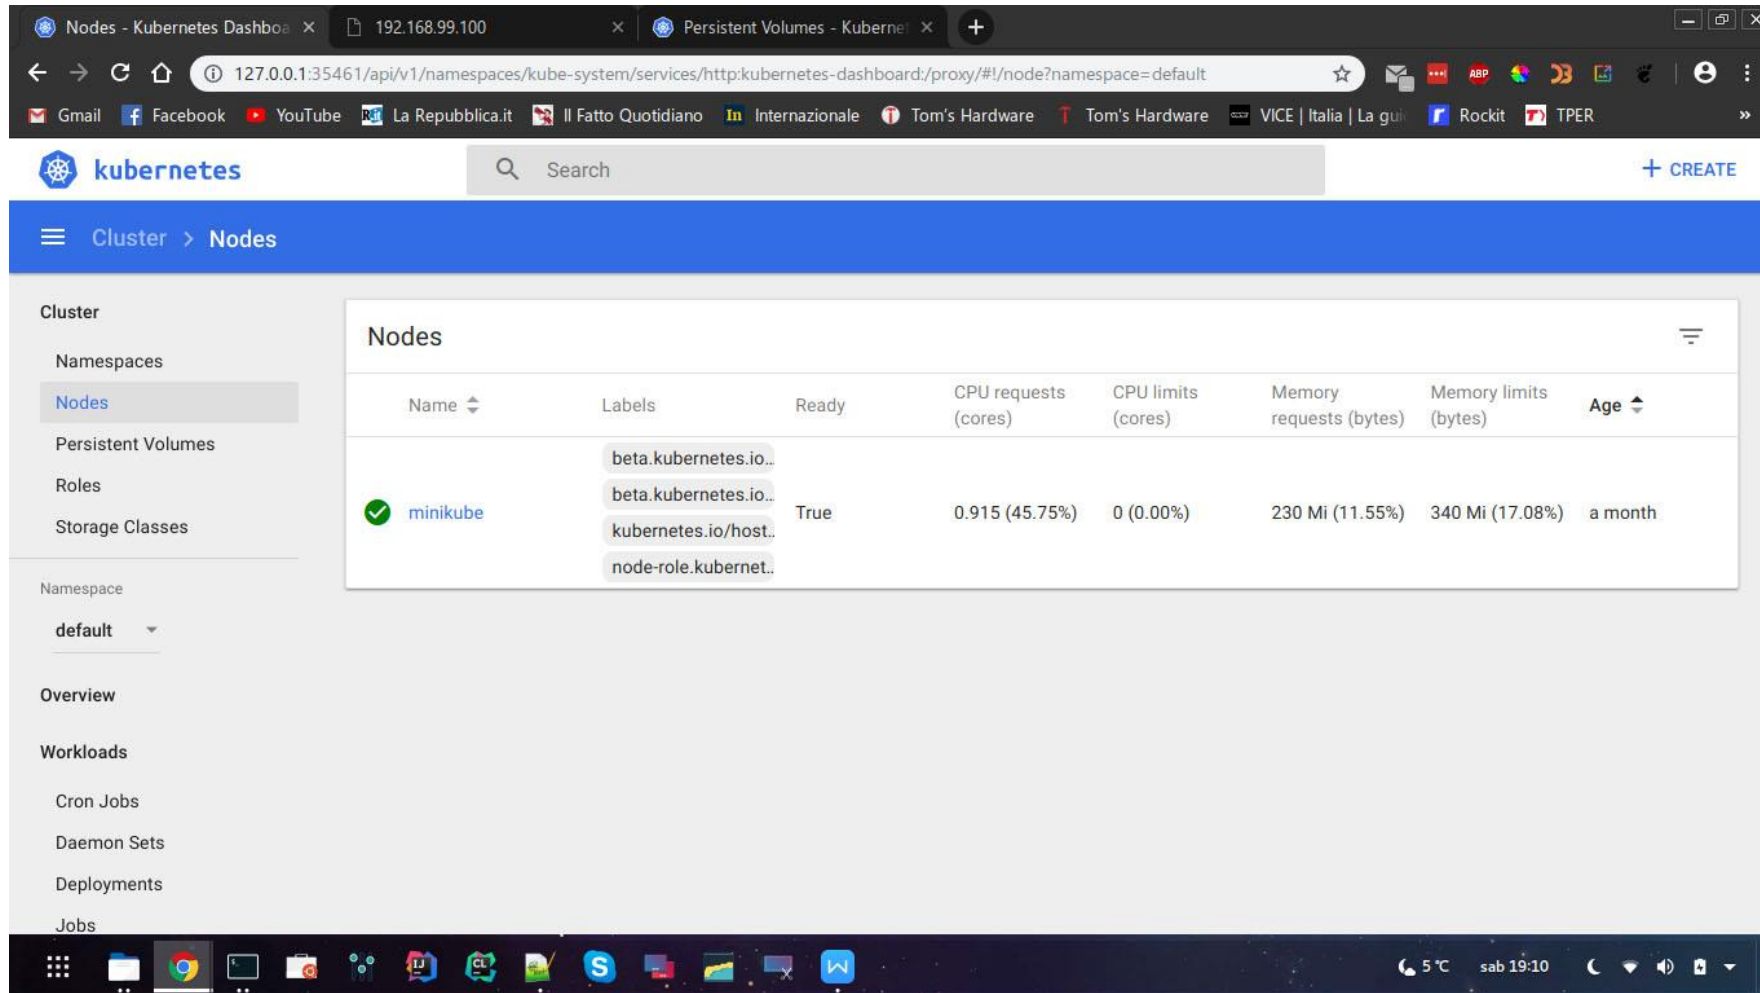
NAME	STATUS	ROLES	AGE	VERSION
minikube	Ready	master	40d	v.1.10.0

## Launch Dashboard

```
$ minikube dashboard
```



## Dashboard



The screenshot shows the Kubernetes Dashboard interface. The top navigation bar includes the Kubernetes logo, a search bar, and a '+ CREATE' button. The main content area is titled 'Nodes' and displays a table of nodes. The table has columns for Name, Labels, Ready, CPU requests, CPU limits, Memory requests, Memory limits, and Age. A single node named 'minikube' is listed with a status of 'Ready'.

Name	Labels	Ready	CPU requests (cores)	CPU limits (cores)	Memory requests (bytes)	Memory limits (bytes)	Age
minikube	beta.kubernetes.io... beta.kubernetes.io... kubernetes.io/host... node-role.kubernet...	True	0.915 (45.75%)	0 (0.00%)	230 Mi (11.55%)	340 Mi (17.08%)	a month

## App organization

- One **pod** for Spring.
- One **pod** MySQL + **Volume**.
- One **service** for MySQL to connect the MySQL pod to the Spring client pod.
- One **service** for Spring to expose publicly the whole app.

Two separated pods in order to scale only the Spring pod.

The first step is to write for each app component the related **Deployment file**.

## spring-deployment.yaml

```
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: spring
5    labels:
6      app: spring
7  spec:
8    replicas: 1
9    selector:
10     matchLabels:
11       app: spring
12  template:
13    metadata:
14      labels:
15        app: spring
16    spec:
17      containers:
18      - name: spring
19        image: davidmonnuar/springbootlibrarydemo:1.0
20        env:
21        - name: DATABASE_DRIVER
22          value: com.mysql.jdbc.Driver
23        - name: DATABASE_HOST
24          value: mysqlldb
25        - name: DATABASE_PORT
26          value: "3306"
27        - name: DATABASE_NAME
28          value: db_example
29        - name: DATABASE_USER
30          value: springuser
31        - name: DATABASE_PASSWORD
32          value: mysql2019
33      ports:
34      - containerPort: 8080
```

## spring-service.yaml

```
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: web-service
5    labels:
6      run: web-service
7  spec:
8    type: NodePort
9    ports:
10     - port: 8080
11       protocol: TCP
12    selector:
13      app: spring
```

## mysql-deployment.yaml

```
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: mysqldb
5  spec:
6    ports:
7      - port: 3306
8    selector:
9      app: mysql
10   clusterIP: None
11 ---
12 apiVersion: apps/v1
13 kind: Deployment
14 metadata:
15   name: mysql
16 spec:
17   selector:
18     matchLabels:
19       app: mysql
20   strategy:
21     type: Recreate
22   template:
23     metadata:
24       labels:
25         app: mysql
26     spec:
27       containers:
28         - image: mysql:5.7
29           name: mysql
```

```
30
31   # Use secret in real usage
32   - name: MYSQL_ROOT_PASSWORD
33     value: password
34   - name: MYSQL_DATABASE
35     value: db_example
36   - name: MYSQL_USER
37     value: springuser
38   - name: MYSQL_PASSWORD
39     value: mysql2019
40 ports:
41   - containerPort: 3306
42     name: mysql
43 volumeMounts:
44   - name: mysql-persistent-storage
45     mountPath: /var/lib/mysql
46 volumes:
47   - name: mysql-persistent-storage
48     persistentVolumeClaim:
49       claimName: mysql-pv-claim
```

```
1 kind: PersistentVolume
2 apiVersion: v1
3 metadata:
4   name: mysql-pv-volume
5   labels:
6     type: local
7   spec:
8     storageClassName: manual
9     capacity:
10      storage: 1Gi
11     accessModes:
12      - ReadWriteOnce
13     hostPath:
14       path: "/mnt/data"
15 ---
16 apiVersion: v1
17 kind: PersistentVolumeClaim
18 metadata:
19   name: mysql-pv-claim
20 spec:
21   storageClassName: manual
22   accessModes:
23     - ReadWriteOnce
24   resources:
25     requests:
26       storage: 1Gi
```

## mysql-pv.yaml

## Deployment (minikubeSetup.sh)

```
1 #!/bin/sh
2 kubectl create -f ./mysql/mysql-pv.yaml
3 kubectl create -f ./mysql/mysql-deployment.yaml
4 kubectl create -f ./spring/spring-deployment.yaml
5 kubectl create -f ./spring/spring-service.yaml
6 kubectl get services
7 minikube ip
```

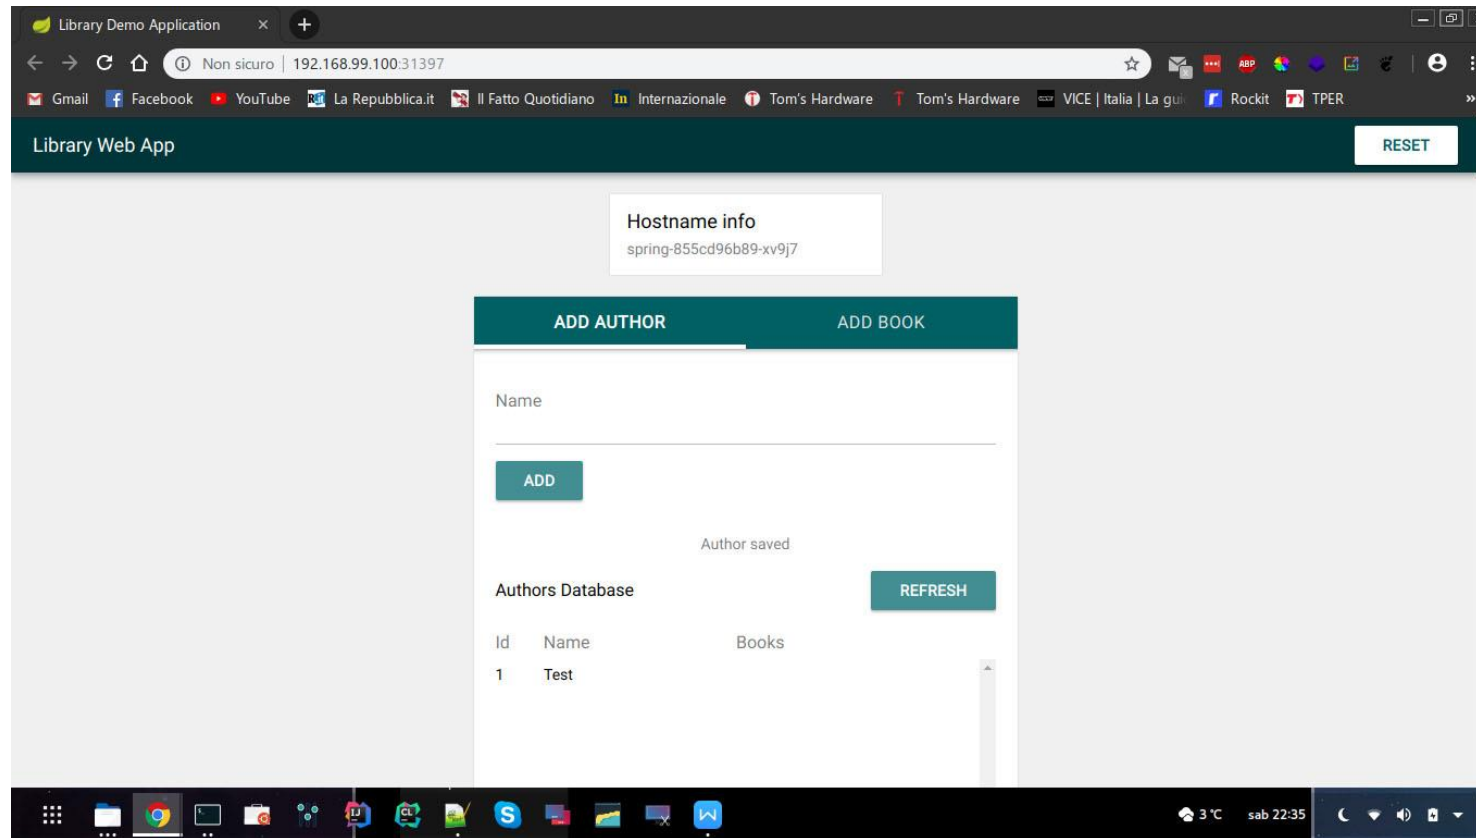
\$ ./minikubeSetup.sh

```
persistentvolume/mysql-pv-volume created
persistentvolumeclaim/mysql-pv-claim created
service/mysqlpdb created
deployment.apps/mysql created
deployment.apps/spring created
service/web-service created
```

NAME	TYPE	CLUSTER_IP	EXTERNAL-IP	PORT(S)	AGE
kubernetes	ClusterIP	10.96.0.1	<none>	443/TCP	47m
mysqlpdb	ClusterIP	None	<none>	3306/TCP	1s
web-service	NodePort	10.96.240.147	<none>	8080:31397/TCP	0s
		192.168.99.100			

To access the deployed app

```
$ minikube service web-service
```



The string reported in the *Hostname Info* box is the pod id in which the container is running.

In Docker that string was the container id.

## Static horizontal scaling by using CLI

```
$ kubectl scale deployments/spring --replicas=2
```

```
$ kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
mysql-7c9fc47c5d-z892k	1/1	Running	0	2h
spring-855cd96b89-2hmwh	1/1	Running	0	2h
spring-855cd96b89-xv9j7	1/1	Running	0	2h

## Dynamic horizontal scaling by defining a scaling policy

- Addon Minikube **metric-server**
- Kubernetes object **Horizontal Pod Autoscaler** (HPA)
- Some changes needed for the spring deployment file

```
...
12   template:
13     metadata:
14       labels:
15         app: spring
16     spec:
17       containers:
18       - name: spring
19         image: davidmonnuar/springbootlibrarydemo:1.0
20         resources:
21           requests:
22             cpu: "500m"
23           limits:
24             cpu: "0.5"
25         env:
26         - name: DATABASE_DRIVER
...
```

As **scaling metric** we focus on **CPU resources usage**

```
$ minikube addons enable metrics-server
```

```
$ kubectl autoscale deployment spring --cpu-percent=50 --min=1 --max=10
```

```
$ kubectl run -i --tty load-generator --image=busybox /bin/sh
#!/ while true; do wget -q -O- http://192.168.99.100:31580/; done
```

## Pod replicas before the workload

```
$ kubectl get hpa
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
spring	Deployment/spring	0%/50%	1	10	1	1m

## Pod replicas 1 minute after the workload start

```
$ kubectl get hpa
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
spring	Deployment/spring	31%/50%	1	10	1	2m

## Pod replicas 3 minutes after the workload start

```
$ kubectl get hpa
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
spring	Deployment/spring	30%/50%	1	10	2	4m

## Pod replicas after the workload end

```
$ kubectl get hpa
```

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
spring	Deployment/spring	0%/50%	1	10	1	9m

Kubernetes is a **powerful** and **effective** tool for containerized apps management.

It is significantly **more complex than Docker Swarm** however it is more flexible.

The **high-level abstractions** that Kubernetes provides to deploy an application let a definition of fine-grained policies and as a consequence the app management is simplified.

Kubernetes is much more than what we've seen so far.

The documentation can be found [here](#).



OpenShift is a family of containerization software developed by Red Hat. Its flagship product is the OpenShift Container Platform (OpenShift from now on), a PaaS for private clouds built around **Docker containers** orchestrated and managed by **Kubernetes** on a foundation of **Red Hat Enterprise Linux**.

## Developement-oriented new features

- Deployment from the Dockerfile.
- Deployment of pre-built Docker images.
- **Deployment directly from the source code** (GitHub).

## Kubernetes extensions

- Several new objects
- Project
- Route

A **constraint** is that in order to use OpenShift you need to use Red Hat Enterprise Linux or Red Hat Atomic.

**Minishift** is a tool that helps you run OpenShift locally by running a single-node OpenShift cluster inside a VM. You can try out OpenShift or develop with it, day-to-day, on your local host. Link [here](#).

The **main advantage** of this plaform is to be closer to software developers than Kubernetes.

IBM Cloud (formerly IBM Bluemix) is a suite of cloud computing services that offers both PaaS and Infrastructure as a Service (IaaS). Among the huge amount of services that IBM Cloud makes available there is also IBM Cloud Kubernetes Service.

In order to use the **Kubernetes-as-a-Service** you need to create a trial account on IBM Cloud and in few seconds you have the access to a single node Kubernetes cluster.

The Kubernetes service is very close to the vanilla Kubernetes provided by Google and it does not contain strong customizations like OpenShift.

Surely the **main advantage** of IBM Cloud is that it quickly provides a Kubernetes cluster off-the-shelf and ready to use without spend too much time in cluster configuration.

Latest news: IBM acquired Red Hat.

# Conclusions

---

**Docker** and **Kubernetes** have completely re-invented the way the cloud industry faces the software development, deployment and maintenance.

The idea of exploiting the native capability of operating systems for isolating the resources by providing the support for containers is successful.

The **containerization** has proven to be an innovative technology that has become in a short time, given the growing adoption rate, a **crucial tool in the modern cloud solutions**.

Source code link: <https://github.com/davidMonnuar/projectActivity>

# Thank you for your attention

Containerization workshop

Infrastructures for Cloud Computing and Big Data M

Dmitrij David Padalino Montenero

Academic Year 2018/2019