Istituto Nazionale di Fisica Nucleare

Containers Orchestration

Daniele Spiga

HPC4L - Training. Beirut, Lebanon
August 21th22th October 2020

Outline

Introduction of the problem

Containers Orchestration

- Overview of the major solutions

Wrap-up

INFN
Istituto Nazionale di Fisica Nucleare

Introduction

- We explored how containers help us to easily create applications that are – as the name says – self-contained.
- We discussed microservices and explored a bit the docker-compose

- What we need then? we’d explore how to effectively orchestrate many containers across distributed hosts.
  - container orchestration.

Three current major solutions

Docker swarm - Apache Mesos - Kubernetes

Docker swarm

Docker Swarm is the traditional way of orchestrating containers with Docker. Compared to other methods we'll see later, it is **relatively easy to use**. Its main features are:

- Cluster management **integrated with Docker Engine**: no other software than docker is needed.
- **Decentralized design**: this means that any node in a Docker Swarm can assume any role at runtime.
- **Scaling**: the Swarm manager can **automatically scale up and down services**, adding or removing tasks.
- **Desired state reconciliation**: if something happens to a Swarm cluster (e.g. some containers crash), the Swarm manager will try to reconcile the state of the cluster to its desired state (e.g. bringing up some more containers).

Docker Swarm (cont)

Docker Swarm features, continued:

- **Multi-host networking**: the Swarm manager can handle an overlay network spanning your services.
- **Service discovery**: there is a DNS server embedded in each Swarm. The Swarm manager discovers services and assigns to each of them a unique DNS name.
- **Load balancing**: you can specify how to distribute services among nodes.
- Secure by default: the communication among all nodes in a Swarm cluster is protected by the cryptographic protocol called TLS (Transport Layer Security).
- **Rolling updates**: if anything goes wrong, you can roll-back a task to a previous version of the service.

Where to get started

INFN
Istituto Nazionale di Fisica Nucleare

Hands-on with Docker Swarm:

- https://docs.docker.com/engine/swarm/swarm-tutorial/

You will learn how to:

- Create a Swarm service
- Deploy a load balancer
- Create swarm cluster
- Create swarm service

Apache Mesos

Apache Mesos is an **open-source cluster manager that provides efficient resource isolation and sharing across distributed applications** (frameworks) ensuring automated self-healing and scalability.

Mesos implements a **two-level meta-scheduler** that provides primitives to express a wide variety of scheduling patterns and use cases

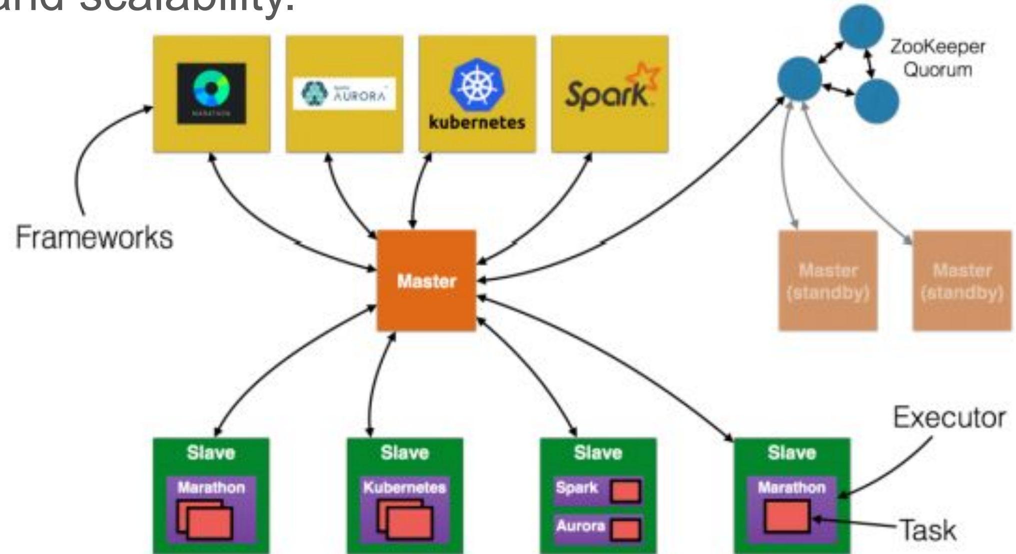

Apache Mesos (cont)

Istituto Nazionale di Fisica Nucleare

Mesos offers a layer of software that organizes the machines (physical servers and/or VMs and/or cloud instances) letting applications draw **from a single pool of intelligently- and dynamically-allocated resources**.

Examples of Mesos frameworks include:

- **Marathon** - a production-grade container orchestration platform designed to launch long-running applications;
- **Chronos** - a distributed fault-tolerant job scheduler; it can be used to run processing tasks.

How it works:

1. Agent 1 reports to the master that it has 4 CPUs and 4 GB of memory free. The master then invokes the allocation policy module, which tells it that framework 1 should be offered all available resources.
2. The master sends a resource offer describing what is available on agent 1 to framework 1.
3. The framework’s scheduler replies to the master with information about two tasks to run on the agent, using <2 CPUs, 1 GB RAM> for the first task, and <1 CPUs, 2 GB RAM> for the second task.
4. Finally, the master sends the tasks to the agent, which allocates appropriate resources to the framework’s executor, which in turn launches the two tasks (depicted with dotted- line borders in the figure). Because 1 CPU and 1 GB of RAM are still unallocated, the allocation module may now offer them to framework 2.

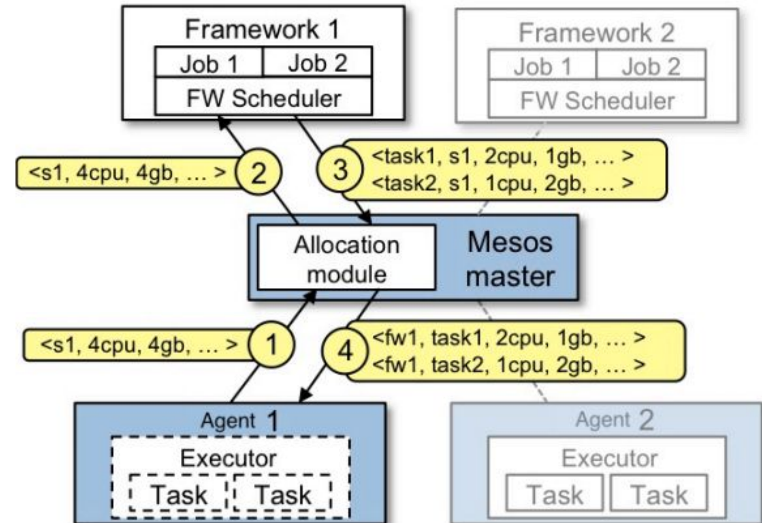

References

http://mesos.apache.org/

https://mesosphere.github.io/marathon/

https://mesos.github.io/chronos/

https://www.consul.io/

Kubernetes

Kubernetes is an **open-source platform that coordinates a highly available cluster of computers that are connected to work as a single unit**. It is backed by Google and RedHat.

- **Applications** need to be **containerized**.
- Kubernetes automates the distribution and scheduling of application containers across a cluster in a fairly efficient way.
- A Kubernetes cluster can be deployed on either physical or virtual machines.

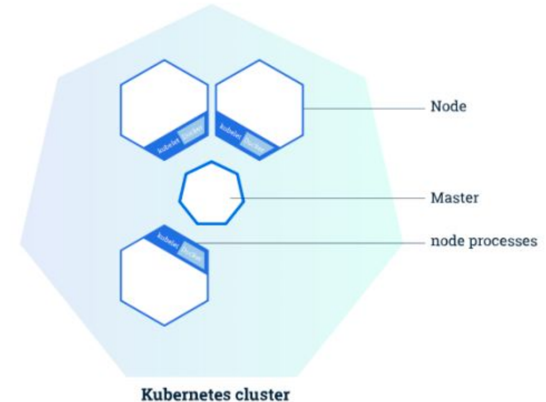

Kubernetes Cluster Resources

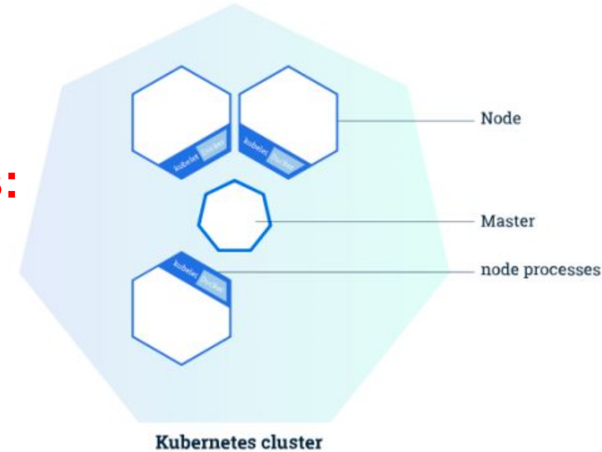

A Kubernetes cluster consists of **two types of resources:**

- **The Master coordinates the cluster**
- **Nodes are the workers that run applications**

The Master is responsible for managing the cluster

- coordinates all activities in your cluster, such as scheduling applications, maintaining applications' desired state, scaling applications, and rolling out new updates.

A node is a VM or a physical computer that serves as a worker machine in a Kubernetes cluster

Kubernetes Master/node processes

Istituto Nazionale di Fisica Nucleare

The Kubernetes Master is a collection of three processes that run on a single node in your cluster, which is designated as the master node. These processes are:

- kube-apiserver
- kube-controller-manager
- Kube-scheduler

Each **individual Node in your cluster runs two processes**:

- kubelet, which communicates with the Kubernetes Master.
- kube-proxy, a network proxy which reflects Kubernetes networking services on each node.

Moreover, **each Node runs a container runtime** (like Docker) responsible for **pulling** the container **image from a registry**, **unpacking** the container, and **running the application.**

- A Kubernetes cluster that handles production traffic should have a minimum of three nodes

Kubernetes Objects

Istituto Nazionale di Fisica Nucleare

Kubernetes contains a number of abstractions that represent the state of your system: deployed containerized **applications and workloads**, their associated **network** and **disk** resources, and other information about what your cluster is doing.

- These abstractions are represented by objects in the Kubernetes API.
- The basic Kubernetes objects include:
  - **Volume**
  - **Namespace**
  - **Deployment**
  - **Pod**
  - **Service**

Kubernetes Deployment

- Once you have a running Kubernetes cluster, you can deploy your containerized applications on top of it. To do so, you create a Kubernetes Deployment configuration.
- The Deployment tells Kubernetes how to create and update instances of your application. Once you've created a Deployment, the Kubernetes master schedules application instances onto individual Nodes in the cluster.
- Once the application instances are created, a Kubernetes Deployment Controller continuously monitors those instances. If the Node hosting an instance goes down r is deleted, the Deployment controller replaces it. This provides a self-healing mechanism to address machine failure or maintenance.
- In a pre-orchestration world, installation scripts would often be used to start applications, but they did not allow recovery from machine failure. By both creating your application instances and keeping them running across Nodes, Kubernetes Deployments provide a fundamentally different approach to applications

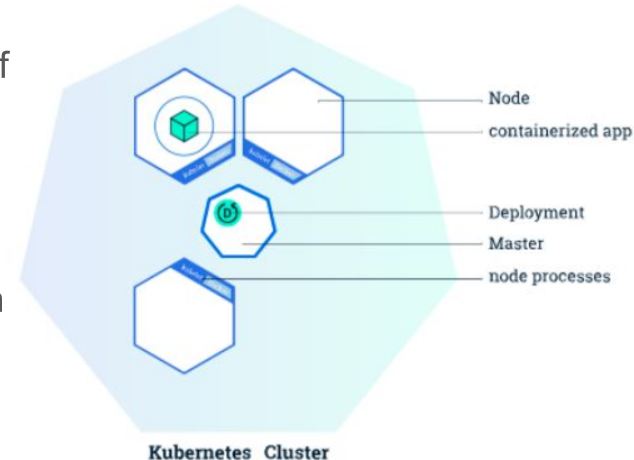

Kubernetes Pod

- A Pod is the basic building block of Kubernetes. It represents a running process on your cluster.
- A Pod encapsulates an application container, storage resources, a unique network IP, and options that govern how the container(s) should run.
- Pods that run a single container. The "one -container - per -Pod" model is the most common Kubernetes use case; in this case, you can think of a Pod as a wrapper around a single container, and Kubernetes manages the Pods rather than the containers directly.
- Pods that run multiple containers that need to work together. A Pod might encapsulate an application composed of multiple co -located containers that are tightly coupled and need to share resources. The Pod wraps these containers and storage resources together as a single manageable entity.

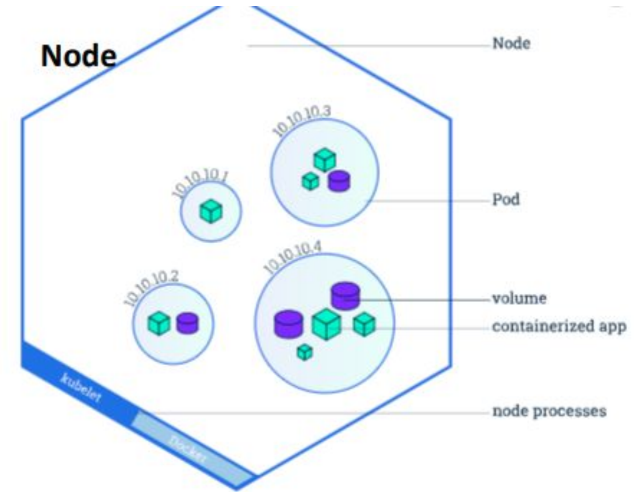

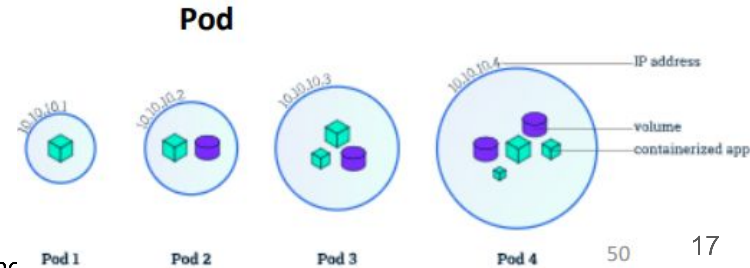

Kubernetes Services

- A Kubernetes Service is an abstraction which defines a logical set of Pods and a policy by which to access it.
- Although each Pod has a unique IP address, those IPs are not exposed outside the cluster without a Service. Services allow your applications to receive traffic.
- Services match a set of Pods using labels and selectors, a grouping primitive that allows logical operation on objects in Kubernetes.
- Labels are key/value pairs attached to objects and can be used in any number of ways:
  - Designate objects for development, test, and production
  - Embed version tags
  - Classify an object using tags

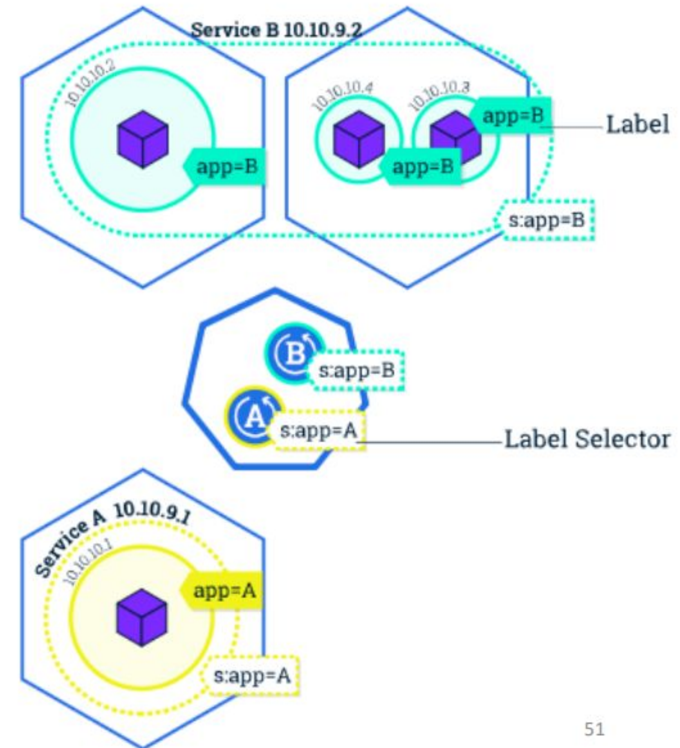

Kubernetes as a Service

Istituto Nazionale di Fisica Nucleare

Deploying and managing a Kubernetes cluster is generally not trivial (that's why Minikube was introduced), since it requires effort and several skills.

- **It would be nice to automatize this part as well, and focus just on deploying our containers on a Kubernetes cluster that somebody else instantiates for us.**

Many Cloud providers give us just that: a Kubernetes as a Service.

- Amazon provides what they call the "Elastic Container Service for Kubernetes", or EKS for short. Other providers have similar offerings.

Docker swarm - Apache Mesos - Kubernetes

INFN Istituto Nazionale di Fisica Nucleare

We have seen (with different degrees of in-depth analysis) **the three current major solutions for container or resource orchestration**.

Question now should be: which one to use ?

Some general considerations on when to use what:

- **Docker Swarm for smaller projects and for testing purposes**. Easy to use if you are already familiar with Docker.
- **For larger, enterprise-like solutions, Kubernetes**. It's also "the Google way of doing it". But mind the rather steep learning curve.
- **Mesos for very large clusters and for workflow-based solutions**. It can be fairly complex, so it might need a sizeable support team.

So let's recap…

Where we are ?

- We have seen Cloud
  - Programming the cloud via PaaS
- We have discussed cloud-native as "synonym" of microservices which are made of container
  - We've discussed dockers
  - We have discussed how combine dockers (docker compose) to build a microservice
- This bring us to the question:
  - What happens when I've "many dockers" and computational resources are distributed ?

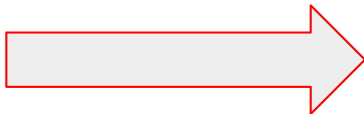

Container Orchestration

So far so good…

Istituto Nazionale di Fisica Nucleare

However, there should be a major question now:

- Where can I deploy / find / access my container orchestrator, more in general where to run my cloud native applications?
- It depends…
    - However I'd generalize the question as: **Where is my infrastructure** ?

We already discussed that you can find Kubernetes as a service (somewhere), that's one option.. A second option is to see **how one could define, and automatize the infrastructure deployment in the Cloud**...