

Kubernetes Tutorial

If you have been working with docker and containers, and want take it to next level ! Well. Kubernetes is an option and in this kubernetes tutorial we shall learn what kubernetes is and what are the use cases that it can cover, with kubernetes concepts and examples.

Kubernetes Tutorial

Kubernetes Tutorial Index

Run Kubernetes on local machine – Kubernetes Tutorial to explain things in detailed on how to setup environment to run kubernetes on your local machine

Kubernetes Overview

Kubernetes is a platform for working with containers, not specifically to docker actually, but containers in general. You may use alternatives to docker to manage the containers, but kubernetes gives you a few key things as a platform that you can build on and extend. To learn about docker, refer [Docker Tutorial](#).

Kubernetes gives you means to do deployments, and is very easy to scale and provides you ability to monitor.

We can enforce desired state management using Kubernetes. Desired State Management meaning we can feed Kubernetes Cluster Services a specific configuration, and Kubernetes takes care to run the configuration in the infrastructure provided.

Kubernetes Cluster

Kubernetes cluster has following Architectural Components

- Kubernetes Master – It has
 - Kubernetes Cluster Services
 - API – Sits in front of all API services
- Kubernetes Worker – Is a container host. It has a
 - Kubelet process – running and is responsible for communicating with Kubernetes Cluster Services.

Kubernetes Cluster consists of one master with Kubernetes Cluster Services running and one or more workers.

Deployment Configuration YAML file

The desired state could be given as a deployment .yaml file. There are many parameters in the cluster that could be given as configuration information. Out these many configurable parameters, there are two fundamental items out of which first is Pod configuration. Pod is smallest unit of deployment in Kubernetes in terms of Kubernetes object model. Which means, in a pod there could be one or more running containers. In order to run the container, an image has to be specified. One or more pod configurations could be provided in the Deployment Configuration YAML file.

Kubernetes Application

Kubernetes contains a master and the master knows about the servers in the cluster. The containers would get deployed in these servers. The process of deployment is simple. Through configuration file, you would tell kubernetes server that what image you would like to create a container from and you would give some criteria. Using the information provided, Kubernetes master creates a deployment, also called application. You may also specify the amount of CPU, amount of RAM and amount of File Storage your application needs in the deployment, and kubernetes will keep track of that for you. Deployment means continuous monitoring, not just initial launching of the container. If your application goes down, kubernetes would know and try through all the possibilities it can to heal the situation. So, it will try to spin another container up and recover for itself.

Kubernetes Scaling

If you are receiving a traffic that is more than expected, and there is a need to scale up your application. The naive way of doing it is increasing the number of servers and deploying application on each of the new servers. But you are always never aware of the amount of traffic and the amount of resources that you need. What kubernetes does is it will figure out what the situation needs and where the required resources are available and hold on to the scales mentioned in the deployment.

At its core, Kubernetes is a platform allowing you to actually maintain deploying containers into production once you get beyond a certain scale. Even if your don't have the need to scale up, you can be benefited from kubernetes with its monitoring capabilities. Kubernetes can help you with automated health checks, rolling restarts and deployments so that you make sure that when you deploy new applications, you are never cutting off anything that needed access to that service.

A Typical Use Case

You provide the Deployment Configuration YAML file to Kubernetes Cluster Services and it is upto Kubernetes Cluster Services to figure out how to schedule the pods in Kubernetes environment and make sure the right number of pods running. If any of the Worker goes down, it is through kubelet process the Kubernetes Cluster Services is informed about the situation. And now if the worker that has gone down is running any of the pods, the environment state does not match the state mentioned in configuration YAML file. Now the scheduler has to make a decision on where to instantiate the abandoned pods.

Conclusion :

In this Kubernetes Tutorial, we have learnt about the concepts of kubernetes with detailed examples.