



kubernetes

Kubernetes: Affinity (Scheduling)

KUBERNETES : Advance of Kuebernetes

- **Kubernetes Scheduler** handle the scheduling of Pods.
- Kubernetes scheduler ensures that the right node is selected by checking the node's capacity for CPU and RAM and comparing it to the Pod's resource requests.
- *Pod affinity* and *pod anti-affinity* allow you to specify rules about how pods should be placed relative to other pods.
- These rules are defined using custom **labels on nodes** and label selectors specified in pods.
- **Pod affinity** can tell the scheduler to locate a new pod on the same node as other pods if the label selector on the new pod matches the label on the current pod.
- **Pod anti-affinity** can prevent the scheduler from locating a new pod on the same node as pods with the same labels if the label selector on the new pod matches the label on the current pod.

KUBERNETES : Advance of Kuebernetes

Requirement of Affinity and Anti-Affinity?

- For a variety of reasons a service or container should only run on a specific type of hardware.
- You want your Pod(s) to end up on a machine with the SSD attached to it.
- You want to co-locate Pods on a particular machine(s) from the same availability zone.
- You want to co-locate a Pod from one Service with a Pod from another service on the same node because these Services strongly depend on each other.

KUBERNETES : Advance of Kuebernetes

- Kubernetes implements node affinity with the **nodeSelector** and **nodeAffinity** fields in PodSpec.
- **nodeSelector** : This is a simple Pod scheduling feature that allows scheduling a Pod onto a node whose labels match the nodeSelector labels specified by the user.
- **Node Affinity** — This is the enhanced version of the nodeSelector introduced in Kubernetes 1.4 in beta. It offers a more expressive syntax for fine-grained control of how Pods are scheduled to specific nodes.
- Label Node in Kubernetes:
`kubectl label nodes <nodeName> key=value`

KUBERNETES : Advance of Kuebernetes

- Some-times we have requirement, not to schedule Pods on specific nodes.
- This feature is the opposite to node and Pod affinity and is known as **node/Pod anti-affinity**.
- **Node Anti-Affinity** — There, we can specify a list of labels which will be compared with the labels of Pods running on the node. If the labels match, the Pod won't be placed alongside with the Pod having this label and running on the node.

Will see you in Next Lecture...

Thank you!

A close-up photograph of a hand holding a black marker, writing the words 'Thank you!' in a cursive script on a white surface. The hand is positioned on the right side of the frame, with the marker tip touching the paper. The background is plain white.

See you in next lecture ...