ONE PAGE SUMMARY

IT & Software › Kubernetes › Pods and Deployments

Introduction to Pods and Deployments

Master Pods and Deployments in Kubernetes for effective application management.

DEFINITION

> Pods are the smallest deployable units in Kubernetes, while Deployments manage the lifecycle of these Pods. Together, they ensure applications run smoothly in a containerized environment.

KEY CONCEPTS

Pod

A Pod is a group of one or more containers that share storage and network resources.

Deployment

A Deployment is a Kubernetes resource that provides declarative updates to Pods and ReplicaSets.

ReplicaSet

A ReplicaSet ensures that a specified number of Pod replicas are running at any given time.

Scaling

Deployments allow you to easily scale the number of Pods up or down based on demand.

Rolling Updates

Deployments support rolling updates, allowing you to update applications without downtime.

EXAMPLES

→ Creating a Pod with a single container running a web server.

→ Using a Deployment to manage multiple replicas of a database service.

→ Scaling a Deployment from 2 to 5 Pods to handle increased traffic.

MEMORY TIPS

★ Remember: Pods are like apartments (containers) in a building (Kubernetes).

★ Think of Deployments as the property manager ensuring everything runs smoothly.

★ Use the acronym PRS: Pod, ReplicaSet, Deployment for key relationships.

COMMON MISTAKES

✗ Confusing Pods with Deployments; they serve different purposes.

✗ Neglecting to specify resource limits for Pods, leading to performance issues.

✗ Overlooking the need for labels and selectors in Deployments.

QUICK RECAP

Pods are the basic units for running containers, while Deployments manage their lifecycle. Understanding how to scale and update Pods through Deployments is crucial for effective application management.