

Kubernetes Workshop

Duration: 1 Day

Description

This workshop aims to provide participants with **hands-on and in-depth knowledge** on **storage management, autoscaling mechanisms, node & pod optimization**, and **advanced scheduling** in Kubernetes environments. Throughout the training, real-world scenarios are used to demonstrate how **persistent storage solutions, scalability approaches**, and **high-availability architectures** are designed and managed on Kubernetes.

What You Will Learn

- The core components of Kubernetes storage architecture
- How to effectively use CSI, PV, PVC, and StorageClass resources
- Managing persistent storage scenarios using different storage drivers
- Applying autoscaling mechanisms in appropriate use cases
- Performing resource optimization at the node and pod levels
- Building high-availability architectures using scheduling, affinity, and anti-affinity
- Running databases on Kubernetes and managing failure scenarios

Audience

- DevOps engineers working with or planning to work with Kubernetes
- Platform and infrastructure teams
- Cloud engineers
- System administrators managing Kubernetes environments
- Technical teams looking to advance their Kubernetes expertise

Outline

CSI, PV, PVC & StorageClass

- Introduction to storage and core concepts
- CSI and its sub-components (Controller, Node Plugin)
- CSI Drivers
 - AWS EBS / EFS

- Azure Disk / File
 - GCE Persistent Disk (PD)
 - Ceph
 - NFS
- PV lifecycle
- Static vs Dynamic provisioning
- Volume types, access modes (RWO, RWX, ROX), and Reclaim policies (Retain, Delete)
- PVC-PV binding process
- PVC resizing
- Using PVCs in Pod and Deployment objects
- VolumeBindingMode types (Immediate vs WaitForFirstConsumer)

Autoscaling Mechanisms

- Introduction to autoscaling
- Horizontal Pod Autoscaler (HPA)
- Vertical Pod Autoscaler (VPA)
- Cluster Autoscaler

Node, Pod & Network Optimization

- Node sizing
- Node taints and tolerations
- CPU & Memory resource limits
- Pod distribution strategies
- Kubernetes networking
- CNI plugins
- Pod-to-Pod and Pod-to-Service traffic

Node Affinity & Scheduling Control

- Scheduling mechanisms

- NodeSelector vs Node Affinity
- Pod Affinity & Anti-Affinity
- High availability using anti-affinity

Database Operation in Kubernetes

- Stateless vs Stateful workloads
- Persistent storage for databases (ReadWriteOnce vs ReadWriteMany)
- Backup & restore strategies
- Pod or node failure simulation

Prerequisites

Docker Fundamentals