ALL DATA ON KUBERNETES

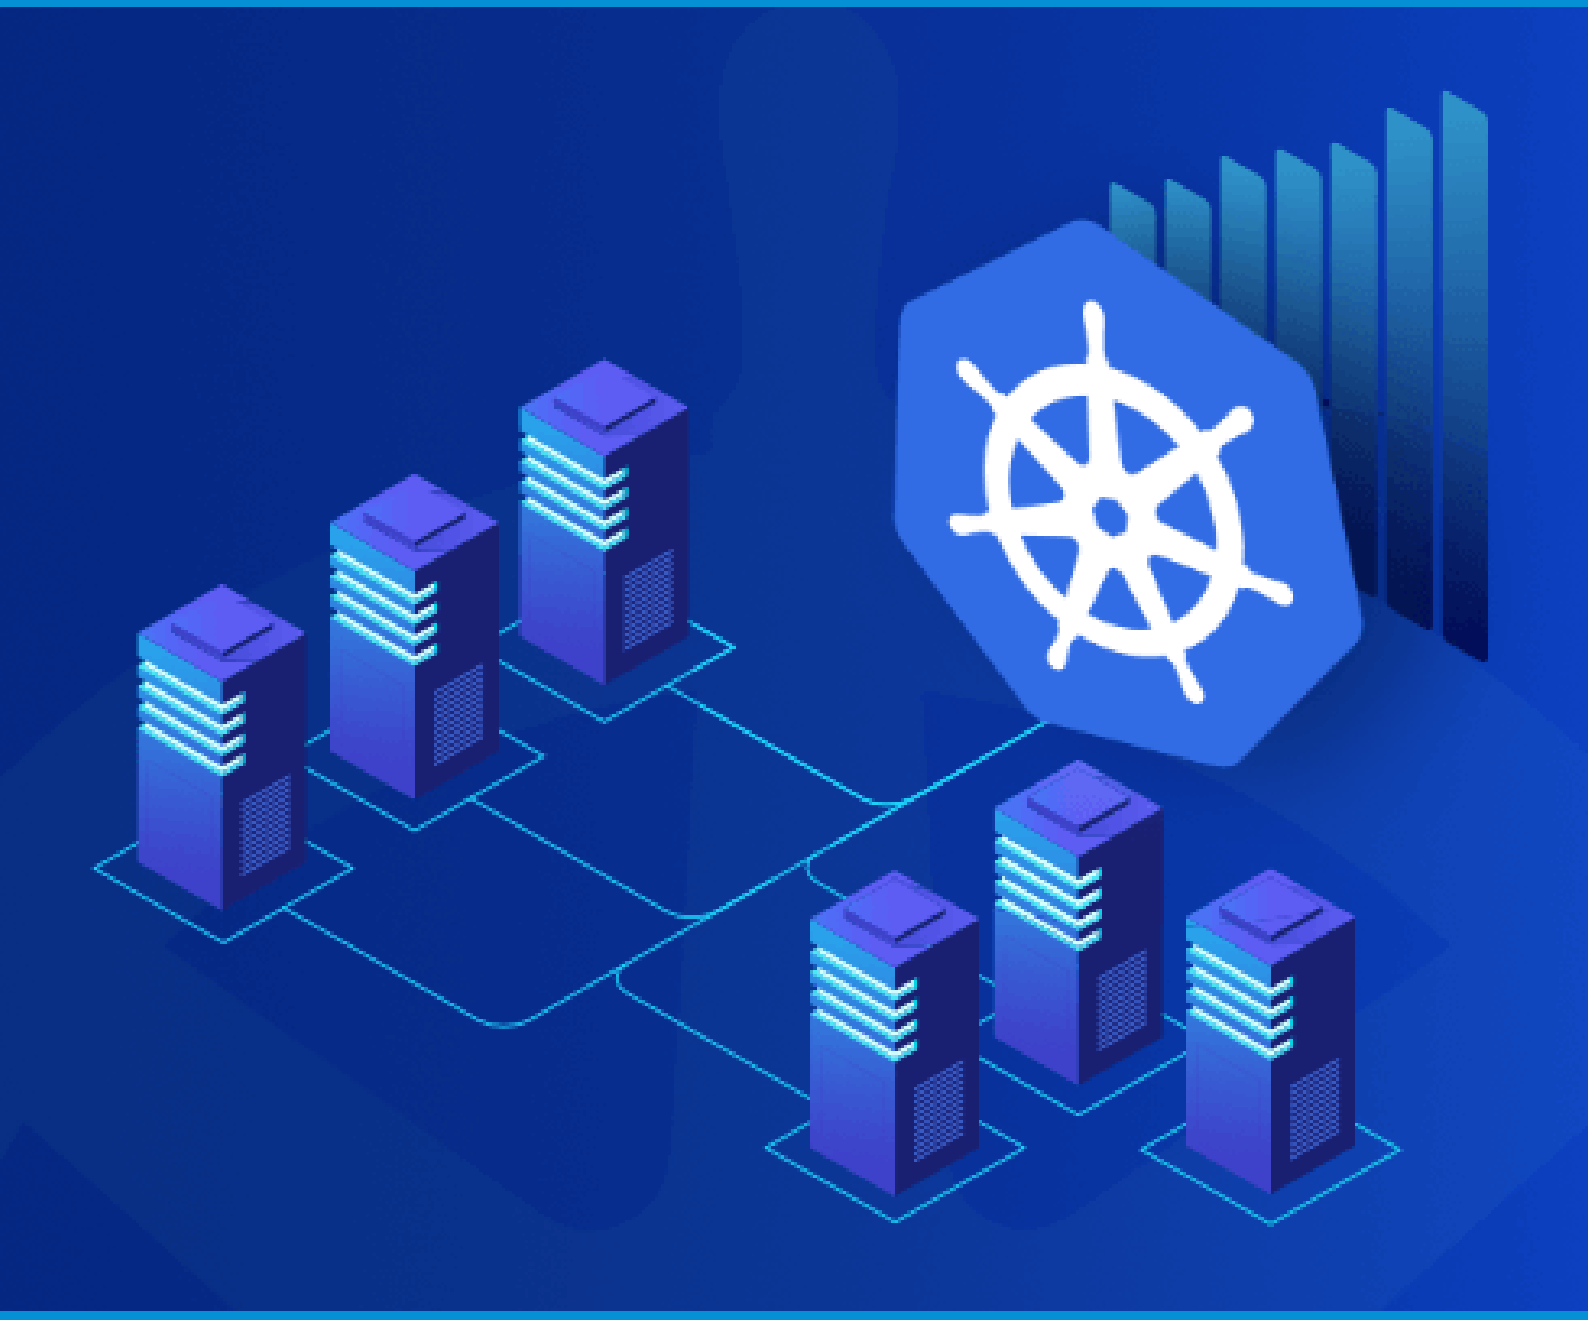

CLOUDAGE

KONDHWA | PUNE | MAHARASHTRA

WWW.CLOUDAGE.CO.IN

COURSE SYLLABUS

All Data on Kubernetes

Module 1: Getting Started

- What is Kubernetes, and Why Do I Need It?
- Installing Kubernetes
- Masters and Nodes

Module 2: Basic Kubernetes Concepts

- Containers and Pods
- Clustering and Nodes
- Networking in Kubernetes
- Kubernetes Architecture and Components

Module 3: Microservices

- What are Microservices?
- Deploying the Robot Shop App
- Deploying a Microservice Application to Kubernetes

Module 4: Understanding Kubernetes Architecture

- Kubernetes Cluster Architecture
- Kubernetes API Primitivesimitives
- Exploring the Kubernetes Cluster via the Command Line

Module 5: Building the Kubernetes Cluster

- Release Binaries, Provisioning, and Types of Clusters
- Installing Kubernetes Master and Nodes
- Building a Highly Available Kubernetes Cluster
- Configuring Secure Cluster Communications
- Running End-to-End Tests on Your Cluster
- Installing and Testing the Components of a Kubernetes Cluster

Module 6: Managing the Kubernetes Cluster

- Upgrading the Kubernetes Cluster
- Operating System Upgrades within a Kubernetes Cluster
- Backing Up and Restoring a Kubernetes Cluster
- Upgrading the Kubernetes Cluster Using kubeadm

Module 7: Cluster Communications

- Pod and Node Networking
- Container Network Interface (CNI)
- Service Networking
- Ingress Rules and Load Balancers
- Cluster DNS
- Creating a Service and Discovering DNS Names in Kubernetes

Module 8: Pod Scheduling within the Kubernetes Cluster

- Configuring the Kubernetes Scheduler
- Running Multiple Schedulers for Multiple Pods
- Scheduling Pods with Resource Limits and Label Selectors
- DaemonSets and Manually Scheduled Pods
- Displaying Scheduler Events
- Scheduling Pods with Taints and Tolerations in Kubernetes

Module 9: Deploying Applications in the Kubernetes Cluster

- Deploying an Application, Rolling Updates, and Rollbacks
- Configuring an Application for High Availability and Scale
- Creating a Self-Healing Application
- Performing a Rolling Update of an Application in Kubernetes

Module 10: Managing Data in the Kubernetes Cluster

- Persistent Volumes
- Volume Access Modes
- Persistent Volume Claims
- Storage Objects
- Applications with Persistent Storage
- Creating Persistent Storage for Pods in Kubernetes

Module 11: Securing the Kubernetes Cluster

- Kubernetes Security Primitives
- Cluster Authentication and Authorization
- Configuring Network Policies
- Creating TLS Certificates
- Secure Images
- Defining Security Contexts
- Securing Persistent Key Value Store
- Creating a ClusterRole to Access a PV in Kubernetes

Module 12: Monitoring Cluster Components

- Monitoring the Cluster Components
- Monitoring the Applications Running within a Cluster
- Managing Cluster Component Logs
- Managing Application Logs
- Monitor and Output Logs to a File in Kubernetes

Module 13: Monitoring Cluster Components

- Troubleshooting Application Failure
- Troubleshooting Control Plane Failure
- Troubleshooting Worker Node Failure
- Troubleshooting Networking
- Repairing Failed Pods in Kubernetes

Google Kubernetes Engine Deep Dive

Module 14: Overview

- Orchestrating Containers with Kubernetes Engine
- Understanding Kubernetes Engine Architecture
- What Does Kubernetes Engine Cost?

Module 15: Creating and Managing Clusters

- Spinning up your First Cluster
- Deploying to your Cluster
- Scaling Clusters and Applications
- Upgrading and Deleting Clusters
- Hands-On Demonstration: Creating and Managing Clusters

Module 16: Handling Security with Kubernetes Engine

- Integrating with Cloud IAM
- Managing Role-Based Access Control
- Establishing a Pod Security Policy
- Implementing Security Protocols
- Hands-on Demonstration: Handling Security with Kubernetes Engine

Module 17: Coordinating Networking and Kubernetes Engine

- Configuring Load Balancers
- Establishing a Network Policy
- Creating a Private Cluster
- Hands-on Demonstration: Coordinating Networking and Kubernetes Engine

Module 18: Expanding Kubernetes Engine

- Implementing Stateful Applications
- Transitioning from On-Prem
- Integrating Other Google Cloud Services
- Monitoring and Logging with Stackdriver

Amazon EKS Deep Dive

Module 19: Exploring EKS

- EKS Architecture
- Configuring an EKS Cluster
- Provisioning Worker Nodes
- IAM Authentication
- Creating an EKS Cluster in AWS

Module 20: Developing for EKS

- Understanding the Application Architecture
- Building from Source
- Publishing to ECR
- Deploying to EKS
- Deploying an Application to EKS

Module 21: EKS in Production

- Autoscaling an EKS Cluster
- Monitoring an EKS Cluster
- Updating EKS in Production
- Autoscaling an EKS Cluster

Module 22: Applying Best Practices

- Logging with CloudTrail
- Continuous Deployment with EKS
- Application Tracing with X-Ray
- Logging to CloudWatch Logs with Fluentd
- Capturing EKS API Calls with CloudTrail

AKS Deep Dive

Module 23: What Is Azure Kubernetes Service?

- What Is Kubernetes?
- Managed Kubernetes

Module 24: Building Your AKS Cluster

- Container Registries and Container Instances
- Service Principal
- Creating and Accessing an AKS Cluster

Module 25: Credentials and Access

- Kubernetes Service Accounts
- Role-Based Access Control (RBAC)

Module 26: Directory Access

- Creating an AKS Cluster with Active Directory Integration
- Authenticating to the AKS Cluster Using Azure Active Directory

Module 27: Network Models

- The Kubenet Network Plugin
- The Azure CNI Network Plugin
- Network Policies

Module 28: Accessing AKS Externally

- Ingress Traffic
- Web Application Firewall (WAF)

Module 29: Storage Types

- Persistent Volumes and Persistent Volume Claims
- Static and Dynamic Volumes

Module 30: AKS Workloads

- Node Sizing
- Azure Site Recovery and Velero

Module 31: Resource Management

- Pod Resource Requests and Limits
- Scaling Options
- Kube-Advisor Tool

Module 32: Advanced Scheduling Features

- Azure Dev Spaces
- Virtual Kubelet

Module 33: Key Metrics

- Azure Monitor for Containers
- Cluster Metrics

Module 34: Collecting Container Logs

- Application Insights
- View Live Logs

About Us

A Big Data consulting and solutions provider offering services and training for Big Data Cloud and Machine Learning

We at CloudAge Provide Hadoop Managed Services to Help Traditional Enterprises adopt Apache Hadoop. we Provide Solutions that includes data preparation, data discovery, data availability, and Data Analytics.

Our expertise with Enterprise Distributions deliver a modern platform for analytics data management offerings, in AWS datacenter, Enterprises get one place to store, access, process, secure, and analyse all their data, empowering them to extend the value of existing investments while enabling fundamental new ways to derive value from their data. Apache Hadoop Open source big data platform is the most widely adopted in the world, and As the leading educator of Hadoop professionals, CloudAge has trained over 1200 individuals worldwide and a seasoned professional services team to help deliver greater time to value.

CloudAge provides best-in-class, technology-managed services and solutions to enterprises that are looking to unlock the potential in their data without the time, cost and complexity associated with traditional big data initiatives. CloudAge delivers an end-to-end solution, so that time is better spent analysing and driving business value from big data. CloudAge provides a full spectrum of services in a private cloud that leverages Hadoop, and helps businesses perform complex analytics and batch-production schedules not possible prior to Hadoop. CloudAge offers a mix of speed, scale, skills, and end-to-end solutions unavailable anywhere else in the big data space.

Insights with Faster Time-to-Value

The analysis of data drives decisions in every business. To gain better business insights, you need to manage the volume, variety, and velocity of data, while applying analytics. With Lenovo-engineered big data validated designs on Lenovo servers, you can harness the power of Apache™ Hadoop® and Apache™ Spark® with Cloudera®, Hortonworks®, IBM® and MapR®. Lenovo servers provide highly reliable and flexible foundations for your business analytics solutions so you can unlock the value of your data and deliver insights faster.

* Outstanding scalability so you can grow as your workloads grow
* Industry-leading transaction processing so you can make better, faster business decisions
* High-throughput capacity that enables you to respond more quickly
* Optimized systems and validated designs for faster time to value