CareerByteCode
Learning Made simple

Kubernetes Bootcamp Syllabus

Day 1 — Kubernetes Fundamentals

- Introduction to Kubernetes
- Containers and Container Orchestration
- Kubernetes Architecture
- Cluster, Node, Control Plane, Worker Node
- Pods and Namespaces
- kubectl Fundamentals
- Deploying Your First Application

Use case:

Deploy a simple Nginx application in a namespace, expose it with a ClusterIP service, and verify access using kubectl.

Day 2 — Kubernetes Networking

- Kubernetes Networking Model
- Pod-to-Pod Communication
- Service Discovery
- Services
    - ClusterIP
    - NodePort
    - LoadBalancer
- DNS in Kubernetes
- Network Policies
- Securing Network Traffic

Use Case:

Deploy two applications, connect one app to the other through a Kubernetes Service, and block unwanted traffic using a basic NetworkPolicy

CareerByteCode
Learning Made simple

Day 3 — Kubernetes Objects and YAML

- Understanding Kubernetes Manifests
- YAML Structure and Best Practices
- Labels, Selectors, and Annotations
- Pods
- ReplicaSets
- Deployments
- DaemonSet
- InitContainers
- Desired State Management
- Managing Resources with kubectl

Use case:

Create a Deployment for a sample web app using YAML, add labels and selectors, and scale it from 2 replicas to 4 replicas.

Day 4 — Deployments and Release Strategies

- Deployment Fundamentals
- Scaling Applications
- Rollouts and Rollbacks
- Deployment Strategies
    - Recreate
    - Rolling Update
    - Blue-Green
- Version Management
- CrashLoopBackOff error and ImagePullBackOff error.
- Zero-Downtime Deployment Concepts

Use case:

Deploy version 1 of an application, upgrade it to version 2 using **Rolling Update**, then switch traffic using a **Blue-Green** style setup.

CareerByteCode
Learning Made simple

Day 5 — Configuration and Storage

- ConfigMaps
- Secrets
- Environment Variables
- Volume Fundamentals
- Persistent Volumes (PV)
- Persistent Volume Claims (PVC)
- Storage Classes
- Stateless vs Stateful Applications

Use case:

Deploy an application that reads configuration from a ConfigMap and database credentials from a Secret, while storing uploaded data in a PVC.

Day 6 — Gateway API

- Introduction to Gateway API
- Gateway API Architecture
- GatewayClass
- Gateway
- HTTPRoute
- Traffic Routing
- Host-Based Routing
- Exposing Applications Externally
- Gateway API Best Practices

Use Case:

Expose two applications under different paths, such as payment.cbc.com and order.cbc.com, using Gateway API instead of Ingress.

CareerByteCode
Learning Made simple

Day 7 — Scheduling and Workload Placement

- Kubernetes Scheduler Fundamentals
- Node Selectors
- Node Affinity
- Node Anti-Affinity
- Pod Affinity
- Pod Anti-Affinity
- Taints
- Tolerations
- Workload Isolation Strategies
- Scheduling Best Practices

use case:

Run a critical application only on dedicated nodes using **taints and tolerations**, and ensure two replicas do not land on the same node using **pod anti-affinity**.

Day 8 — Resource Management and Reliability

- Resource Requests and Limits
- CPU and Memory Management
- Quality of Service (QoS) Classes
- Understanding OOMKilled
- CPU Throttling
- Readiness Probes
- Liveness Probes
- Startup Probes
- Horizontal Pod Autoscaler (HPA)
- Troubleshooting Application Health

Use case:

Run a memory-heavy app with low limits, observe an **OOMKilled** event, then fix it by tuning requests, limits, and probes.

CareerByteCode
Learning Made simple

Day 9 — CI/CD and GitOps for Kubernetes

GitHub Actions CI

- CI/CD Fundamentals
- GitHub Actions Overview
- Building Container Images
- Automated Testing
- Container Registry Integration
- Managing Secrets in Pipelines

Deploying to Kubernetes

- Automated Deployment Workflows
- Deployment Validation
- Release Automation

GitOps with Argo CD

- GitOps Fundamentals
- Argo CD Architecture
- Applications and Projects
- Sync Operations
- Self-Healing and Drift Detection
- Rollbacks and Version Management

Use case:

Push code to GitHub, let **GitHub Actions** build and publish the image, then use **Argo CD** to sync the application into Kubernetes automatically from Git.

CareerByteCode
Learning Made simple

Day 10 — Helm and AI-Powered Kubernetes Operations

Helm Fundamentals

- Why Helm
- Helm Architecture
- Helm Charts
- Chart Structure
- values.yaml
- Installing Applications with Helm
- Upgrading and Rolling Back Releases
- Customizing Deployments Using Values Files

K8sGPT and Ollama

- Introduction to K8sGPT
- Connecting K8sGPT with Ollama
- AI-Assisted Cluster Analysis
- Troubleshooting Common Kubernetes Issues

Practical Troubleshooting Lab

- CrashLoopBackOff
- ImagePullBackOff
- OOMKilled

use case:

Package an application as a **Helm chart**, deploy it with custom values, then use **K8sGPT + Ollama** to analyze a broken pod and explain the issue in simple terms.

Final Capstone Project

Project Theme

Deploy and stabilize a customer-facing static web application using GitHub Actions, Argo CD, Gateway API, Helm, and AI-assisted troubleshooting with K8sGPT.

Real-world problem to solve

- the app is not coming up after deployment,
- one pod keeps restarting,

- traffic is not reaching the service correctly,
- or the new release causes unstable behavior.