Kubernetes Fundamentals

Course KUBFND 2 Days Instructor-led, Hands on

Introduction

In the future, organizations that succeed will be those that put data at the forefront of their business strategy. This Data Literacy training is designed to give you a high level overview of the key topics in Data Science and Machine Learning. Designed exclusively for students who want to learn about the basics of data science and machine learning at a high level, without needing to learn how to code or cover complex mathematics. In this course you'll learn the fundamental concepts relating to data, allowing you to understand what makes data suitable for data analysis, visualization and machine learning. Then we'll give you a quick overview of important statistical topics, such as mean, standard deviation, and the normal distribution. Afterwards you will learn the different ways data scientists are able to visualize data to convey their ideas in a clear manner. We'll also teach you about the machine learning process, acquiring data, cleaning data, and an overview of the train/test split philosophy that supervised learning adheres to. Then we'll show you some examples of regression and classification algorithms, as well as how to evaluate their results. We'll also explore what the future holds by taking a peek at the bleeding edge of AI and ML, including DALLE-2 and GPT-3!

Who Should Attend?

- Anyone who plans to work with Kubernetes at any level or tier of involvement
- Any company or individual who wants to advance their knowledge of the cloud environment
- Any Developer/Engineer/SA beginning their Kubernetes Learning
- Project Managers/Leads

Course Objectives

By the end of this course, you will have a solid understanding of:

- Understand Kubernetes architecture and core cluster components
- Deploy and manage workloads using Kubernetes primitives (Pods, ReplicaSets, Deployments)
- Create, modify, and manage Kubernetes resources using YAML manifests
- Configure container health checks, security contexts, and observability tools
- Manage compute resources through limits, requests, and quotas
- Collect and analyze logs from containers and Pods
- Configure and use persistent storage with PersistentVolumes and PersistentVolumeClaims
- Implement and scale application workloads (Deployments, ReplicaSets, DaemonSets, StatefulSets)

Contact ISInc for more information at 916.920.1700 or by visiting our website at http://www.isinc.com

- Expose applications internally and externally using Services (ClusterIP, NodePort, LoadBalancer)

Prerequisites

Prior to attending, students should be familiar with the following:

- Linux for Absolute Beginners
- Or any other Formal Linux Training

Course Outline

Module 1: From Containers to Kubernetes

- Kubernetes Architecture
- Pods and the Control Plane
- Kubernetes the Alta3 Way
    - Deploy Kubernetes using Ansible

Module 2: Cluster Basics

- Namespaces and Fundamental Kubectl Commands
    - Creating and Configuring Namespaces
- Kubectl get and sorting
    - Listing Resources with kubectl get

Module 3: Pod Basics

- YAML
- Manifests for Pods
    - Create and Configure Basic Pods
- API Versioning and Deprecations

Module 4: Container Health, Security, and Observability

- Kubectl port-forward
- Kubectl exec and cp
    - Performing Commands inside a Pod
- Readiness and Liveness Probes
    - Implement Probes and Health Checks
- Pod Security Contexts
    - Applying Security Contexts

Module 5: Resource Management

- Limits, Requests, and Namespace ResourceQuotas
    - Defining Resource Requirements, Limits and Quotas
    - Kubectl Top and Application Monitoring

Contact ISInc for more information at 916.920.1700 or by visiting our website at http://www.isinc.com

- Admission Controller
    - Create a LimitRange AdmissionController

Module 6: Logging

- Utilize Container Logs
    - Kubectl Log Command

Module 7: Persistent Storage

- Persistent Volumes, Claims, and StorageClasses
    - Using PersistentVolumeClaims for Storage

Module 8: Deployments

- Labels
    - Labels and Selectors
- Annotations
- ReplicaSets
    - Create and Configure a ReplicaSet
- Deployments - Purpose and Advantages
    - Create and Configure a Deployment

Module 9: Services: ClusterIP, NodePort, Loadbalancer

- Networking with Services
    - Expose Applications via Services

Module 10: EXTRA: Ephemeral Storage

- ConfigMaps and Volume Mounting
    - Persistent Configuration with ConfigMaps
- Secrets
    - Create and Consume Secrets

Module 11: EXTRA: Persistent Storage

- Persistent Volumes with CSI
    - CSI Storage Solution: NFS

Module 12: EXTRA: Helm

- Helm
    - Making Charts and Templates with Helm
    - Deploy Existing Packages via Helm