

CELEBRATING
14 YEARS

QualityThought[®]
Transforming Dreams! Redefining Future!



Khaja Ibrahim
Sr. Cloud Architect
Trained 10,000+ Students & Professionals



A gateway to your
Bright Future in the
IT Industry

📞 95151 51992, 99637 99240



Lets Start Here

Hello there! It's a pleasure to connect with you all, my future DevOps engineers. I'm Khaja, and I'm absolutely thrilled to have the opportunity to guide you on this exciting journey into the world of DevOps. Over the years, I've had the privilege of teaching thousands of students, and I must say, watching them embark on their DevOps adventure has been nothing short of amazing.



Khaja Ibrahim
Sr. Cloud Architect
Trained 10,000+ Students & Professionals



Our journey together will be full of hands-on experiences, practical exercises, and real-world projects that will not only build your technical skills but also prepare you to excel in the fast-paced, ever evolving world of Cloud And I wanted to make it more individual for you based on your background from which you are

Transition into DevOps

- ▶ zero knowledge or very little IT knowledge
- ▶ system administrator
- ▶ software developer
- ▶ test engineer
- ▶ network engineer
- ▶ database administrators
- ▶ middle ware administrators

DevOps is not just a methodology; it's a cultural shift, a mindset that bridges the gap between development and operations to faster collaboration, automation, and continuous improvement. This course is designed to be your gateway to the dynamic and evolving landscape of DevOps, equipping you with the skills and insights needed to thrive in a world where speed, efficiency, and innovation are paramount.

In an era where the pace of technological change is unprecedented, DevOps emerges as the linchpin that accelerates the delivery of high-quality software. It's about breaking down silos, automating repetitive tasks, and fostering a culture of collaboration and continuous learning. Whether you're a seasoned IT professional or someone venturing into the tech realm, the principles of DevOps are applicable and transformative.

In this course, we will delve into the core principles and practices of DevOps, exploring topics such as version control, continuous integration, continuous delivery, and infrastructure as code. But more than just theory, we'll engage in hands-on exercises that simulate real-world scenarios, ensuring you gain practical experience in implementing DevOps methodologies.

Our goal is not just to equip you with a toolkit of DevOps tools (although we will certainly do that!), but to instill in you the DevOps mindset. You'll learn not just how to use tools like Jenkins, Docker, Kubernetes and Ansible, but also how to integrate them into a seamless and automated pipeline that enhances collaboration and accelerates delivery.

Your success in this course is not just a milestone; it's a testament to the power of embracing DevOps principles. As you progress through the modules, you'll find that DevOps is not just a set of practices; it's a philosophy that transforms the way we approach software development and delivery.


So, buckle up for a journey where you'll not only acquire technical skills but also cultivate a mindset that embraces change, values collaboration, and drives continuous improvement. Whether you're looking to advance in your current role, transition into a DevOps position, or simply stay ahead of industry trends, DevOps is your compass to navigate the ever-evolving landscape of technology.

Welcome to the DevOps revolution. Let's build, automate, and innovate together!

Software Development Concepts




As a DevOps engineer you may not be programming the application, but as you are working closely with the development team to improve and automate tasks for them,



How developers work and collaborate
(Agile, AWS DevOps workflows)

How applications are configured
(Build & packaging Tools)



Linux For the Job

Linux is the foundation of most server environments, and DevOps and cloud engineers need to be proficient in Linux to manage and automate server configurations.

Basic Concepts

- ▶ Shell Commands
- ▶ File System & Permissions
- ▶ User Management
- ▶ SSH Key Management
- ▶ Package Management
- ▶ Process Management
- ▶ Disk Management
- ▶ Networking
- ▶ Shell Scripting and Automation
- ▶ Network Configurations
- ▶ DNS, Load Balancers and Proxies

Ansible

Why Ansible?

Ansible is a powerful open-source automation tool used in software development and IT operations for a variety of reasons, especially when building complex systems.

Think of Ansible as a super-efficient butler for your complex system. Here's why you'd want this helpful butler:

Automation: Ansible can take care of repetitive and time-consuming tasks automatically, like setting up servers, installing software, and making sure everything runs smoothly. This frees up your time and reduces the chances of mistakes.

Consistency: Imagine having a butler who makes sure everything in your house is in the right place and working perfectly every day. Ansible does the same for your system, ensuring everything is set up exactly as it should be, no matter how many parts your system has.

Organization: Ansible helps you keep track of all the things happening in your system, like making sure the lights turn on before the coffee maker starts. It helps you manage complex processes, so nothing gets missed.

Scalability: When you have a big house, you need your butler to manage everything smoothly, whether you have a few guests or a big party. Ansible does the same for your system, making sure it works just as well whether you have a few users or thousands.



Safety and Security:

Ansible also acts like a security guard, making sure all the doors are locked and the alarm is set. It helps you keep your system safe from threats and ensures it follows all the rules.

Ease of Use:

Just like a butler who knows your preferences, Ansible can be customized to fit your system's unique needs. It's like having a butler who knows exactly how you like things done.

Ansible is like having a reliable and efficient helper that takes care of all the complex tasks in your system, making sure everything runs smoothly and consistently while giving you more time to focus on the important stuff.

What you would be learning?

- ▶ Inventory and Ad-Hoc Commands
- ▶ Roles and Role-Based Organization
- ▶ Error Handling and Handlers
- ▶ Ansible Vault and Security
- ▶ Ansible Best Practices and Optimization
- ▶ Integrating Ansible with Other Tools
- ▶ Real-World Complex System Scenarios
- ▶ Troubleshooting and Debugging
- ▶ Case Studies

Additional Tasks

Multi-Tier Web Application: Build an Ansible playbook to deploy a multi-tier web application, including setting up web servers, application servers, and databases.

Automated Deployment Pipeline: Set up a CI/CD pipeline using Jenkins, GitLab CI/CD, or Travis CI, where Ansible automates the deployment of code to different environments (e.g. development, staging, production).

Server Configuration Management: Develop Ansible roles and playbooks to manage server configurations, including software installations, system updates, and security hardening.

Application Deployment: Develop Ansible playbooks for deploying various applications (e.g., web apps, databases) with different configurations.

Rolling Updates: Create Ansible scripts to perform rolling updates of applications to minimize downtime.

Version Control System using Git

Why Git?

Git is essential for DevOps engineers working on complex systems because it plays a crucial role in enabling collaboration, automation, and version control within the DevOps lifecycle. Here's why Git is important for DevOps engineers in simple terms

Collaboration Backbone:

Git serves as the backbone for collaboration among DevOps engineers, developers, and other team members. It allows multiple individuals to work on the same codebase simultaneously without conflicts, ensuring that everyone is on the same page.

Version Control Magic:

Git is like a magic time machine for code. It lets DevOps engineers track changes to code over time, making it easy to revert to earlier versions if something goes wrong. This ensures that the system can be rolled back to a stable state in case of issues.

Infrastructure as Code (IaC): DevOps often involves managing infrastructure as code, where server configurations and deployments are defined in code. Git helps DevOps engineers version and manage this code, making it reproducible, shareable, and auditable.

Automated Pipelines:

Git integrates seamlessly with CI/CD (Continuous Integration/Continuous Deployment) pipelines. DevOps engineers can use Git to trigger automated builds, tests, and deployments whenever changes are pushed to the code repository, streamlining the software delivery process.

Branching Strategies: Git's branching and merging capabilities enable DevOps engineers to implement advanced strategies, such as feature branching and release management. This makes it possible to work on new features or fixes separately from the main codebase and merge changes in a controlled manner.

Configuration Management: Git can be used to manage configuration files for various infrastructure components. DevOps engineers can version control configurations, making it easier to track changes, audit configurations, and maintain consistency across environments.

Collaborative Problem Solving: In complex systems, issues and bugs are inevitable. Git provides tools for issue tracking and collaborative problem-solving. DevOps engineers can use Git alongside issue tracking systems to manage and resolve problems efficiently.

Cross-Team Collaboration: DevOps often involves collaboration between different teams, such as development, operations, and security. Git serves as a common platform for these teams to coordinate their efforts, ensuring that everyone has access to the latest code and configurations.



Auditing and Compliance:

Git's version history and commit logs provide an audit trail of changes made to the system. This is valuable for compliance purposes and for demonstrating that security and regulatory requirements are met.

Git is like the control center for DevOps engineers. It helps them manage code, configurations, and collaboration seamlessly, ensuring that complex systems can be developed, deployed, and maintained with efficiency, reliability, and accountability. It's a foundational tool in the DevOps toolbox.

What you would be learning?

- ▶ Introduction to Version Control and Git
- ▶ Basic Git Commands
- ▶ Branching and Merging
- ▶ Remote Repositories and GitHub
- ▶ Collaborative Workflows
- ▶ Git Rebase and Interactive Rebase
- ▶ Git Cherry-Pick and Reset
- ▶ Git Hooks and Continuous Integration
- ▶ Branching Strategies

Static Code Analysis using SonarQube

Why SonarQube ?

Imagine SonarQube as a guardian for your code, helping DevOps engineers ensure that the software they build is of high quality and free from hidden issues. Here's why DevOps engineers need SonarQube in simple terms

Code Quality Detective: SonarQube acts like a detective for your code, carefully inspecting it for problems, like spelling mistakes in a book. It identifies code quality issues, security vulnerabilities, and bugs that might otherwise go unnoticed.

Early Warning System: Just as a smoke detector alerts you to a fire before it spreads, SonarQube warns DevOps engineers about code problems early in the development process. This prevents issues from becoming major fires (or bugs) later on.

Security Guard: SonarQube plays the role of a security guard, scanning your code for weaknesses and vulnerabilities that could be exploited by malicious actors. It helps ensure that your software is as secure as possible.

Consistency Checker: Like a proofreader checking for consistent language usage in a document, SonarQube enforces coding standards and best practices across your codebase. This helps maintain a uniform and readable codebase, making it easier for teams to collaborate.

Continuous Improvement Coach: SonarQube doesn't just find issues; it also offers suggestions for improvement. It's like having a coach who provides tips and guidance on how to make your code better and more efficient.

Tracking Progress: SonarQube keeps track of your code's health over time. It's similar to tracking your fitness progress in a workout app. DevOps engineers can see if code quality is improving or if there are recurring issues that need attention.

Team Collaboration: SonarQube acts as a common ground for developers, testers, and DevOps engineers to discuss and prioritize code improvements. It helps teams work together to maintain and enhance the quality of the software.

In simple terms, SonarQube is your code's guardian angel, watching over it, finding issues, and helping DevOps engineers ensure that the complex systems they build are reliable, secure, and of the highest quality. It's an essential tool for building and maintaining software that meets high standards.

What you would be learning?

- ▶ SonarQube Basics
- ▶ Code Analysis Concepts
- ▶ Quality Gates and Metrics
- ▶ Identifying Code Issues
- ▶ SonarQube Plugins
- ▶ Integration with CI/CD Pipelines
- ▶ Custom Rules and Quality Profiles
- ▶ Reporting and Notifications

Artifactory (Jfrog)

Why Artifactory?

Think of Artifactory as a treasure chest for DevOps engineers, storing valuable software artifacts like code, libraries, and dependencies. Here's why DevOps engineers need Artifactory in simple terms

Artifact Safekeeper:

Artifactory is like a secure vault where you can store and manage all the pieces of software that make up your applications. This includes code, libraries, and other important building blocks.

Version Control:

Just as a library organizes books by category and edition, Artifactory helps DevOps engineers keep track of different versions of software artifacts. This ensures that the right versions are used in the right places.



Efficient Sharing:

Artifactory acts like a library card system. It allows DevOps engineers to share and distribute software artifacts easily among team members, ensuring everyone has access to the same resources.

Dependency Management:

When building complex systems, software often relies on other pieces of software. Artifactory helps manage these dependencies, making sure the right components are available when needed.

Reproducible Builds:

Artifactory ensures that builds can be recreated exactly as they were before, just like following a recipe. This reproducibility is crucial for troubleshooting and maintaining consistent software.

Security Guard:

In a world full of threats, Artifactory acts as a security guard. It scans and verifies software artifacts to prevent vulnerabilities and unauthorized access.

High Availability:

Artifactory is always available, like a 24/7 library. It ensures that software artifacts are accessible when developers and systems need them, without interruptions.

Remote Repositories:

Just as you can borrow books from libraries in different cities, Artifactory can connect to remote repositories. This means you can access software artifacts from all over the world.

Metadata Magic:

Artifactory provides metadata about artifacts, like a catalog for books. This metadata helps DevOps engineers search, organize, and understand the software they're using.

In simple terms, Artifactory is like a well-organized and secure library for software artifacts. It keeps everything in order, ensures that software is safe and accessible, and makes it easy for DevOps engineers to build and maintain complex systems with confidence.

What you would be learning?

- ▶ Artifactory Basics
- ▶ Managing Artifacts
- ▶ Repository Management
- ▶ Version Control and Dependency Management
- ▶ Metadata and Properties
- ▶ Integration with Build Tools
- ▶ CI/CD Pipeline Integration

Azure DevOps

Why Azure DevOps?

Think of Azure DevOps as a toolbox that helps DevOps engineers build, test, and deliver software with speed, quality, and collaboration. Here's why DevOps engineers need Azure DevOps in simple terms:

Project Organizer:

Azure DevOps acts like a project manager's whiteboard. It helps DevOps engineers plan and organize their software projects, breaking them into smaller tasks and tracking progress.

Code Workshop:

Just as a workshop needs tools, Azure DevOps provides a space for DevOps engineers to work on code collaboratively. It stores and manages code, ensuring everyone is on the same page.

Testing Playground:

Imagine a playground where kids test their ideas. Azure DevOps offers tools for automated testing, making sure that the software works correctly and reliably.

Delivery Pipeline:

Like a conveyor belt in a factory, Azure DevOps automates the process of packaging and delivering software. It ensures that the right version of the software reaches the right place at the right time.

Collaboration Hub:

DevOps is all about teamwork. Azure DevOps serves as a virtual meeting room where developers, testers, and operations folks can work together, share insights, and resolve issues.

Quality Inspector:

Just as a food inspector ensures that restaurants follow health standards, Azure DevOps checks the quality of code. It spots issues and helps fix them before they become problems.

Feedback Collector:

Like a suggestion box, Azure DevOps gathers feedback from users and team members. This feedback helps DevOps engineers improve the software continuously.

Security Guard:

In a world full of cyber threats, Azure DevOps acts as a security guard. It ensures that the software is safe from vulnerabilities and follows security best practices. Reports and Dashboards: Azure DevOps provides reports and dashboards like a scoreboard in a game. It helps DevOps engineers see how the project is performing and where improvements are needed.

Integration Hub:

In a complex system, various tools need to work together. Azure DevOps acts as an integration hub, connecting with other tools and services that DevOps engineers use.

In simple terms, Azure DevOps is like a Swiss Army knife for DevOps engineers. It offers a set of essential tools and a collaborative workspace that helps streamline the software development and delivery process, ensuring that the complex systems they build are efficient, secure, and reliable.

What you would be learning?

- ▶ Azure DevOps Overview
- ▶ Setting up a Sample Project
- ▶ Version Control with Azure Repos
- ▶ Azure Pipelines Basics
- ▶ Advanced Pipeline Configurations
- ▶ Test Planning with Azure Test Plans
- ▶ Automated Testing with Azure DevOps
- ▶ Application Insights and Monitoring
- ▶ Continuous Improvement and Feedback
- ▶ Release gates, deployment strategies, and infrastructure as code (IaC) integration

Additional Tasks

- ▶ CI/CD Pipeline for Web Application
- ▶ Infrastructure as Code (IaC) with Azure DevOps
- ▶ Microservices Deployment
- ▶ Release Management for Multi-Tier Application
- ▶ Infrastructure Compliance and Security Scanning
- ▶ Deployment to Multiple Cloud Providers
- ▶ High Availability (HA) Deployment
- ▶ Serverless Application Deployment



GitHub Actions

GitHub Actions is a CI/CD (Continuous Integration and Continuous Deployment) tool integrated into GitHub that allows you to automate, test, and deploy code directly from your repositories. It enables developers to create workflows that are triggered by events such as code pushes, pull requests, or issue creation.

1. Integration with GitHub

- ▶ **Seamless CI/CD:** GitHub Actions is built directly into GitHub, so it integrates smoothly with GitHub repositories. This allows you to trigger workflows based on repository events like pushes, pull requests, releases, or issues.

2. Flexibility and Customization

- ▶ **Custom Workflows:** GitHub Actions allows you to define custom workflows using YAML files, which makes it highly customizable. You can run tests, build projects, deploy code, and automate almost anything based on triggers.

3. Scalability

- ▶ GitHub Actions supports running jobs on different operating systems (Linux, macOS, Windows), and you can configure multiple jobs to run in parallel or sequentially, providing robust scalability.

4. Rich Ecosystem

- ▶ **GitHub Marketplace:** There's a vast marketplace of pre-built actions that cover a wide range of use cases, like running unit tests, deploying to cloud providers (AWS, Azure, Google Cloud), and automating release processes.

5. Cost-Effective

- ▶ GitHub provides free minutes for public repositories and a generous amount of free minutes for private repositories, which can make it a cost-effective solution compared to other CI/CD services.

6. Built-in Security Features

- ▶ **Secrets Management:** GitHub Actions has built-in support for securely storing and accessing secrets, like API keys or credentials.

7. Self-Hosted Runners

- ▶ If you need to run workflows on your own infrastructure, GitHub Actions provides support for self-hosted runners. This gives you full control over the environment and hardware where your workflows run.

8. Support for Matrix Builds

- ▶ GitHub Actions can run tests on multiple configurations (e.g., different versions of languages, operating systems, etc.) using matrix builds. This is particularly useful for ensuring compatibility across environments.

9. Event-Driven Automation

- ▶ Besides CI/CD, GitHub Actions can be used to automate various other tasks like issue management, labeling, updating documentation, or automating repetitive development tasks based on GitHub events.

Use Cases for GitHub Actions:

- ▶ **CI/CD pipelines:** Build, test, and deploy applications automatically.
 - ▶ **Automation of repetitive tasks:** Automate project management tasks such as labeling issues, triaging pull requests, or notifying team members.
 - ▶ **Infrastructure as Code (IaC):** Deploy and manage cloud infrastructure using IaC tools like Terraform.
- Open-source projects: Automatically test and validate code contributions from the community.

What you would be learning?

- ▶ Module 1 : Introduction to GitHub Actions
- ▶ Module 2: Getting Started with GitHub Actions
- ▶ Module 3: Events that Trigger Workflows
- ▶ Module 4: Jobs and Runners
- ▶ Module 5: Actions
- ▶ Module 6: Contexts and Expressions
- ▶ Module 7: Environment Variables and Secrets
- ▶ Module 8: Advanced Workflow Features
- ▶ Module 9: Testing and Debugging Workflows
- ▶ Module 10: Security and Compliance
- ▶ Module 11 : Integration with External Tools and Services
- ▶ Module 12: Real-World Use Cases and Projects

GitOps with ARGO-CD

GitOps with Argo CD is a powerful and efficient approach to managing and deploying applications in Kubernetes environments, especially for those looking to implement Continuous Deployment (CD) practices. Here are some key reasons why GitOps with Argo CD is a great choice

Why GitOps with ARGO-CD

1. Git as the Single Source of Truth

- ▶ **Declarative Infrastructure:** GitOps uses Git repositories to store the desired state of your application and infrastructure. Every change is version-controlled, meaning Git serves as the single source of truth for your system.

2. Continuous Deployment

- ▶ Argo CD continuously monitors your Git repositories for changes to the desired state. When changes are detected (like a new version of an application or updated configuration), Argo CD automatically applies those changes to your Kubernetes cluster.

3. Declarative and GitOps-Native Approach

- ▶ **Kubernetes-Native:** Argo CD is designed specifically for Kubernetes, making it a great fit for managing applications in a Kubernetes environment. It applies the declarative configuration principles of Kubernetes to application management.

4. Ease of Rollbacks and Auditing

- ▶ **Git Rollbacks:** Rollbacks in GitOps with Argo CD are as simple as reverting a Git commit. Once the commit is reverted, Argo CD will detect the change and apply it to your cluster.

5. Self-Healing & Drift Detection

- ▶ **Automated Drift Correction:** Argo CD continuously compares the live state of the cluster with the desired state in Git. If it detects any drift (manual changes or misconfigurations), it can alert users or automatically sync the cluster back to the desired state.

6. Better Developer Productivity

- ▶ **Focus on Code, Not Infrastructure:** Developers work with familiar Git workflows (committing, pushing, pull requests) without worrying about Kubernetes-specific tools or manual deployment steps.

7. Security & Policy Enforcement

- ▶ **Approval Gates & Policies:** You can configure Argo CD to require manual approval for specific environments (like production) while allowing automatic deployments for lower environments (like staging or development).

8. Scalability and Multi-Cluster Support

- ▶ Argo CD can manage deployments across multiple Kubernetes clusters, making it a great tool for organizations with complex environments.

9. Visualization and Monitoring

- ▶ **Real-Time Monitoring:** Argo CD provides a web UI and CLI for visualizing the current state of your applications. You can see the state of your Kubernetes cluster, identify drift, and monitor deployments in real-time.

10. Integrates Well with Other Tools

- ▶ **CI/CD Pipelines:** Argo CD can integrate with existing CI pipelines to trigger deployments based on successful builds. For example, CI tools like Jenkins, GitHub Actions, or CircleCI can push changes to Git, and Argo CD will take care of deploying those changes.

What you would be learning?

- ▶ Module 1 : Introduction to GitOps
- ▶ Module 2: Getting Started with Argo CD
- ▶ Module 3: Working with Applications in Argo CD
- ▶ Module 4: GitOps Workflows with Argo CD
- ▶ Module 5: Advanced Argo CD Features
- ▶ Module 6: Sync and Application Management
- ▶ Module 7: Integrating Argo CD with Other Tools
- ▶ Module 8: GitOps Best Practices and Patterns
- ▶ Module 9: Scaling GitOps in the Enterprise
- ▶ Module 10: Troubleshooting and Debugging
- ▶ Module 11 : Hands-On Projects and Labs

Terraform

Why Terraform?

Think of Terraform as a magical blueprint for building and managing the infrastructure that supports your software applications. Here's why DevOps engineers need Terraform in simple terms

Infrastructure Wizardry:

Terraform is like a wizard's spellbook for creating and managing servers, databases, and other infrastructure components. It allows DevOps engineers to describe their desired infrastructure in a simple, human-readable language.

Consistency Enforcer:

Just as a recipe ensures that you make the same delicious dish every time, Terraform ensures that your infrastructure is consistent. It creates and configures resources exactly as you specify, reducing errors and surprises.

Efficiency Booster:


Terraform automates the process of creating, modifying, and destroying infrastructure. It's like having an army of helpers who can set up servers and services in minutes instead of hours or days.

Multi-Cloud Harmony:

Terraform is cloud-agnostic, meaning it works with different cloud providers like AWS, Azure, and Google Cloud. It lets DevOps engineers manage infrastructure across multiple clouds with a single set of commands.

Version Control Friend:

Just as you save different versions of your document, Terraform lets you version-control your infrastructure. This means you can track changes over time and easily roll back to previous configurations if needed.



Collaboration Facilitator:

Terraform enables teamwork. Multiple DevOps engineers can work on the same infrastructure code, and Terraform helps merge their changes and maintain consistency.

Risk Minimizer:

Like a safety net, Terraform can help recover from disasters. If something goes wrong, you can use your Terraform code to rebuild your infrastructure exactly as it was before.

Security Sentinel:

Terraform helps ensure that your infrastructure is configured securely. It can enforce security policies and best practices, reducing vulnerabilities.

In simple terms, Terraform is like having a magic wand for setting up and managing the servers, databases, and other technology that your software needs to run. It makes sure everything is built correctly, saves time, and keeps everything organized and consistent, whether you're working in one cloud or many.

What you would be learning?

- ▶ Terraform Basics
- ▶ Terraform Configuration Language (HCL)
- ▶ Resource Types and Providers
- ▶ Variables and Data Sources
- ▶ Terraform State
- ▶ Collaboration and Version Control
- ▶ Modules and Code Organization
- ▶ Workspaces and Environments

Additional Tasks

- ▶ Multi-Cloud Deployment
- ▶ Web Application Deployment
- ▶ Container Orchestration
- ▶ Infrastructure for Microservices
- ▶ High Availability (HA) Architecture
- ▶ Multi-Region Deployment
- ▶ Continuous Integration and Deployment (CI/CD) Pipeline
- ▶ Monitoring and Alerting Infrastructure
- ▶ Logging and Observability Stack

Docker

Why Docker ?

Think of Docker as a magic box where you can package everything your software needs to run, no matter where it goes. Here's why DevOps engineers need Docker in simple terms

Consistency Creator:

Docker ensures that your software runs the same way on every computer, from your laptop to a server in the cloud. It's like a recipe that guarantees the same delicious meal every time.

Easy Shipping:

Just as you can ship a package anywhere in the world, Docker lets you package your software and all its parts (code, libraries, settings) into a neat container. This container can be sent to any computer that runs Docker, and it will work without any fuss.

Resource Saver:

Docker helps save resources because it uses less space and memory compared to traditional methods. It's like packing your suitcase efficiently for a trip, making the most of the available space.

Quick Start:

With Docker, you can start new projects or services in minutes, rather than hours or days. It's like having a chef who prepares everything in advance, so you can cook a meal faster.

Scaling Made Simple: When your software becomes popular, Docker makes it easy to run more copies of it to handle the extra load. It's like having a bakery that can quickly make more loaves of bread when there are more customers.

Security Shield: Docker helps keep your software safe. It's like having a protective shield around your application, making sure it doesn't interfere with other software on the same computer.

Flexibility Friend:

Docker plays well with others. It can work with many different types of software and technologies, making it flexible and adaptable to your needs.

In simple terms, Docker is like a magic container that holds everything your software needs, making it easy to move around, quick to start, and consistent in its behavior, while saving resources and keeping your software secure. It's a valuable tool for DevOps engineers to manage and deploy complex systems.



What you would be learning?

- ▶ Docker Basics
- ▶ Working with Containers
- ▶ Docker CLI and Docker Hub
- ▶ Container Networking
- ▶ Data Management with Volumes
- ▶ Docker Compose
- ▶ Docker Images and Registries

Additional Tasks

- ▶ Dockerized Web Application
- ▶ Microservices Deployment
- ▶ Containerized Dev/Test Environments
- ▶ Dockerized Legacy Application
- ▶ Container Security Scanning
- ▶ Custom Docker Images and Repositories
- ▶ Docker as a CI/CD Runner

Kubernetes

Why Kubernetes?

Think of Kubernetes as a smart traffic cop for your software. Here's why DevOps engineers need Kubernetes in simple terms:

Traffic Manager:

Kubernetes manages the flow of your software, just like a traffic cop directs cars on the road. It ensures that your applications run smoothly, traffic is balanced, and there are no jams.

Auto-Pilot for Containers:

Imagine having a team of robots to handle containers (the boxes that hold your software). Kubernetes automates tasks like starting, stopping, and moving containers, making your software run reliably.

Scaling Guru:

When your software becomes popular, Kubernetes can quickly create more copies to handle the load, like opening extra lanes on a highway during rush hour.

Fault Detective:

If something goes wrong with a container, Kubernetes can replace it with a fresh one, just like getting a new car if your old one breaks down.



Multi-Cloud Navigator:

Kubernetes can work on different cloud platforms (like AWS, Azure, or Google Cloud), allowing you to switch between them or use them all together.

App Upgrades without Traffic Jams:

You can update your software without causing disruptions. Kubernetes directs traffic away from old versions and gradually shifts it to the new one, like smoothly rerouting traffic during road construction.

App Safety Net:

Kubernetes ensures your apps are always available. If a server crashes, it moves your containers to a healthy one, like quickly rerouting traffic if there's an accident on the road.

App Uniformity:

It makes sure all your containers are set up the same way, like making sure all cars on the road follow the same rules.

In simple terms, Kubernetes is like a super-smart traffic controller for your software, ensuring it runs smoothly, adapts to changes, and stays available, no matter where it's deployed. It's a valuable tool for DevOps engineers to manage complex systems.

What you would be learning?

- ▶ Kubernetes Basics
- ▶ Pods and Containers
- ▶ Deployments and Services
- ▶ Config Maps and Secrets
- ▶ Persistent Storage
- ▶ Scaling Applications
- ▶ Load Balancing and Ingress
- ▶ Monitoring and Logging
- ▶ Helm and Package Management
- ▶ Azure Kubernetes Service
- ▶ AWS Elastic Kubernetes Service

Additional Tasks

- ▶ Microservices Deployment
- ▶ Multi-Environment Deployment
- ▶ Stateful Applications
- ▶ Continuous Deployment Pipeline
- ▶ Ingress Controller and Load Balancing
- ▶ Kubernetes Monitoring and Logging Stack
- ▶ GitOps Workflow
- ▶ Kubernetes Deployment with Jenkins
- ▶ Kubernetes Deployment with AzureDevOps

Observability using Prometheus & Grafana

Observability using Prometheus is essential in modern software systems, especially for microservices and cloud-native architectures, due to its robust features designed for monitoring and alerting. Here's why Prometheus is a popular choice for observability:

Why Observability using Prometheus and Grafana

1. Metrics-Based Monitoring

- ▶ Prometheus is a powerful, open-source metrics-based monitoring system that provides real-time insights into your applications and infrastructure. Unlike traditional logging, which focuses on event data, Prometheus emphasizes metrics, allowing you to:

2. Pull-Based Data Collection

- ▶ Prometheus operates on a pull-based model, meaning it scrapes targets (like application services) for metrics at defined intervals. This allows for:

3. Time-Series Data Storage

- ▶ Prometheus stores time-series data, making it well-suited for observability. Time-series data includes:

4. Efficient Querying with PromQL

- ▶ Prometheus provides PromQL (Prometheus Query Language), a flexible and powerful language for querying and aggregating time-series data. With PromQL, you can:
 1. Create detailed queries to analyze system performance.
 2. Perform complex calculations like rates, averages, and percentiles over time.
 3. Visualize data in dashboards or set up alerts based on specific metrics.

5. Alerting Capabilities

- ▶ Prometheus integrates with the Alertmanager component to provide alerting functionality. Alerts can be triggered based on PromQL queries, and the Alertmanager handles:
 1. Routing alerts to different notification systems (e.g., email, Slack, PagerDuty).
 2. Grouping alerts to reduce noise and manage alert fatigue.
 3. Escalation policies to ensure critical alerts reach the right teams.

6. Designed for Cloud-Native and Microservices Environments

- ▶ Prometheus was designed with cloud-native and microservices architectures in mind, which makes it an excellent choice for modern distributed systems:
- ▶ Service Discovery: Prometheus can automatically discover services running in environments like Kubernetes, AWS, or Docker Swarm, thanks to its built-in service discovery mechanisms.
- ▶ Scaling: It can monitor hundreds or thousands of microservices efficiently, pulling data from multiple sources and scaling with your system.
- ▶ Ephemeral and dynamic environments: Its ability to scrape metrics from dynamic sources (e.g., containers and pods) makes it well-suited for environments with short-lived workloads.

7. Integration with Ecosystem Tools

- ▶ Prometheus integrates with a wide variety of tools for visualization, alerting, and data storage. Notable integrations include:
 1. Grafana: For creating rich and interactive dashboards.
 2. Thanos: For long-term storage of metrics and scaling Prometheus across large clusters.
 3. Alertmanager: For centralized alert handling and routing.
 4. OpenTelemetry: Prometheus can work alongside OpenTelemetry for broader observability, integrating with tracing and logging data.

8. Open-Source and Community Support

- ▶ As an open-source tool, Prometheus has a strong community and wide adoption, especially in the CNCF (Cloud Native Computing Foundation) ecosystem. This means:
 1. Frequent updates and improvements.
 2. Access to community-created exporters and integrations.
 3. Reduced costs compared to proprietary monitoring tools.

9. Multi-Dimensional Data Model

- ▶ Prometheus uses a label-based data model, which allows for multi-dimensional queries. This model is flexible and powerful for dissecting metrics based on attributes, such as:
 1. Instance (e.g., server1, server2)
 2. Environment (e.g., dev, prod)
 3. Service (e.g., frontend, backend)
- ▶ This multi-dimensional labeling enables users to segment, analyze, and understand data from various sources in a more meaningful way.

10. High Availability and Fault Tolerance

- ▶ Prometheus can be configured to run in HA (High Availability) mode, where multiple Prometheus instances scrape the same targets. Additionally, Prometheus can work with tools like Thanos or Cortex for long-term storage and high availability across multiple clusters.

What do we learn?

- ▶ When learning observability using Prometheus and Grafana, you gain a deeper understanding of system monitoring, performance management, and the principles of modern DevOps. Here's a summary of the key learnings:

1. Core Concepts of Observability
2. Metrics Collection and Storage (Prometheus)
3. Monitoring System Health
4. Data Visualization (Grafana)
5. Alerting and Incident Management
6. Performance Tuning and Optimization
7. Infrastructure as Code (IaC)
8. Scalability and Reliability
9. Kubernetes Monitoring
10. Extensibility and Ecosystem

This knowledge is vital for anyone working in DevOps, SRE (Site Reliability Engineering), cloud infrastructure, or system administration, ensuring systems stay healthy and performant in production environments.

Python For DevOps

► Python is widely used in DevOps for its simplicity, flexibility, and the power it brings to automate processes, manage infrastructure, and streamline CI/CD pipelines. Below are the key areas where Python is applied in DevOps:

1. **Core Python**
2. **Command line Applications**
3. **Deploying Python Applications to cloud**
4. **ORM Frameworks (SQLAlchemy)**
5. **Building an API's**
6. **Cloud and API Integration**

Python simplifies interaction with cloud APIs, automating processes such as scaling instances, provisioning resources, and managing user data.

Common Libraries:

- 6.1. **Boto3 (AWS)**
- 6.2. **Google Cloud SDK**
- 6.3. **Azure SDK for Python**

Conclusion

Python is an indispensable tool in the DevOps ecosystem. Its versatility allows DevOps engineers to automate tasks, manage infrastructure, integrate with CI/CD pipelines, handle container orchestration, and more. By leveraging Python, DevOps teams can increase productivity, improve efficiency, and streamline the entire development and deployment process.