Trainer: **Raghu K**

Azure DevSecOps Training
Complete Syllabus 2026

Master Cloud-Native DevOps, Security, and AI-Powered Automation with an industry-aligned, hands-on curriculum that takes you from fundamentals to production-grade expertise.

14 Weeks | 5 Days / Week | 100% Hands-On | Live + Lab

13 Modules + Capstone | CKA / AZ-104 / Terraform Cert Aligned

Trainer: **Raghu K** | www.learndevopsonline.com

WHY CHOOSE US

Why This Training?

Built by practitioners, refined by thousands of students, and designed for real-world impact.

Industry-Aligned Curriculum

Inspired by and aligned with **roadmap.sh/devops**, our syllabus covers every skill modern DevOps engineers need — from Linux to AI-powered operations. No gaps, no fluff.

One Trainer, All Topics

Raghu K teaches every single module. No hand-offs between instructors means a **consistent, coherent learning experience** where every concept builds on the last.

Live Enterprise-Grade Project

You won't just learn tools in isolation. You'll build and deploy a **complete microservices application** on Azure with full CI/CD, security gates, and observability — exactly like a production environment.

8500+ Students Trained

A proven track record with over **8,500 students** successfully trained and placed in top companies. Our graduates are working at startups and Fortune 500 enterprises alike.

"If you find content more comprehensive than ours at this price, let us know and take the course for FREE!"
— That's how confident we are in this syllabus.

YOUR TOOLKIT

Tools You'll Master

Gain hands-on proficiency with the most in-demand DevSecOps tools in the industry.

Linux OS	**Bash** SCRIPTING	**Git** VCS	**GitHub Actions** CI/CD	**Azure Cloud** CLOUD
Python SCRIPTING	**Ansible** CONFIG MGMT	**Terraform** IAC	**Docker** CONTAINERS	**Kubernetes** ORCHESTRATION
Helm PACKAGE MGR	**ArgoCD** GITOPS	**SonarQube** CODE QUALITY	**Vault** SECRETS	**Trivy** SECURITY
Prometheus MONITORING	**Grafana** DASHBOARDS	**ELK Stack** LOGGING	**OpenTelemetry** OBSERVABILITY	**New Relic** APM
Istio SERVICE MESH	**AI / AIOps** AUTOMATION			

CONTENTS

Table of Contents

14 weeks of structured, progressive learning across 13 modules.

DETAILED SYLLABUS

Course Modules

Deep-dive into every module. Each topic is covered with live demos and hands-on labs.

MODULE 01

Linux & Shell Scripting

Week 1

1.1 Linux Fundamentals

- ✓ Anatomy of Linux
- ✓ Linux Command Line Introduction
- ✓ Directory Structure in Linux

1.2 File & Directory Management

- ✓ ls, touch, rm, cp, mv commands
- ✓ cd, pwd, mkdir, rm -r, cp -r commands
- ✓ Rename directories using mv

1.3 Text Processing & Search

- ✓ vim editor
- ✓ grep (word search), awk (column filter), head & tail (line filters)
- ✓ find (find files)

1.4 System Utilities

- ✓ curl (download content)
- ✓ tar and unzip (extract compressed files)
- ✓ Pipes in the command line

1.5 Process & User Management

- ✓ ps and kill (process management)
- ✓ useradd (user management)
- ✓ sudo command & sudoers configuration

1.6 Software & Service Management

- ✓ dnf (package management)
- ✓ systemctl (service management)

1.7 Permissions & Networking

- ✓ chown, chgrp (ownership)
- ✓ chmod (file permissions)
- ✓ Network info commands

1.8 Shell Scripting

- Problems solved by shell scripting, use cases and limitations
- Types of shell scripting & why Bash is famous
- Writing, getting, and executing shell scripts
- Variables: declaring, accessing, command substitution, arithmetic substitution
- Special variables, taking input, single vs double quotes, exit status
- Functions: declaring, scope, return status
- Conditional statements: if, case
- Loops: while, for
- Input redirectors, SED editor, sourcing files

1.8 Networking for DevOps NEW

- TCP vs UDP, ports, sockets
- DNS resolution — how domains resolve to IPs
- HTTP/HTTPS, status codes (2xx, 3xx, 4xx, 5xx)
- SSL/TLS handshake & certificates
- CIDR notation & subnetting basics
- Load balancing concepts (L4 vs L7)
- Firewalls, iptables, network troubleshooting (ping, traceroute, netstat, curl, dig, nslookup)

MODULE 02

Azure Cloud Fundamentals

Week 2

Note: Advanced Azure services (AKS, Key Vault, WAF, FinOps) are covered contextually throughout the course in their respective modules.

2.1 Getting Started

- Azure Account SignUp
- Resource Groups
- Azure Regions and Zones

2.2 Networking & Security

- Public IP & Private IP
- Network Security Groups on Subnets and Resources
- NSG Inbound & Outbound Rules
- Virtual Network (VNET)
- Subnets and Private Subnets
- Route tables Private subnets
- NAT Gateway and Implementation
- Subnet Categorization
- Subnet Network Security groups
- Azure VNET Peering
- Azure Landing Zone Overview

2.3 Compute (Virtual Machines)

- Virtual Machines
- Creation of a Virtual Machine
- Connect to Linux Azure VM Server from Windows / Mac / Linux Desktop
- Azure VM Spot Instances
- Azure VM Spot Eviction Policies

2.4 DNS

- Azure DNS Zones
- Azure Private DNS Zones

2.5 Identity & Access (Microsoft Entra ID)

- Microsoft Entra ID (Azure Active Directory [AAD])
- Users and groups in Entra ID
- Role assignments to Users and Groups
- Azure Role assignments and Role definitions
- App Registrations (Service Principals)

2.6 Storage

- Storage Accounts
- Azure Blob Storage
- Azure Cool Storage
- Azure Archive Storage
- Encrypting Blob storage
- Versioning data in Blob storage

2.7 Security & Encryption

- Azure Key Vault (Keys, Secrets & Certificates)
- Data encryption using Keys (Azure Key Vault)
- Azure WAF
- Azure DDoS Protection
- App Service Certificate

2.8 Database Services

- CosmosDB Overview, Create CosmosDB Cluster, CosmosDB in private connectivity
- Azure Cache for Redis Overview, Create Redis Cluster, Redis in private connectivity
- Azure Database for MySQL, Create MySQL Server, MySQL Server in private connectivity
- Backup Policies

2.9 Load Balancing & Auto Scaling

- Azure Application Gateway
- Azure Load Balancer
- Public & Internal Load Balancers
- Backend Pools
- Expose private apps running on private servers over Load Balancer
- Azure Autoscale
- Virtual Machine Scale Sets

2.10 Container Services

- Azure Container Registry
- Azure Kubernetes Service
- Azure managed identities (System & User managed)

2.11 Monitoring & Management

- Azure Portal
- Log Analytics
- Application Insights
- Azure Monitor
- Azure Resource Manager (ARM Templates)

2.12 Other Essential Services

- Azure Functions
- Azure CDN
- Azure management groups
- Azure Native ISV Service
- Azure compute gallery

2.13 FinOps & Cost Optimization NEW

- Azure Cost Management & Billing
- Resource tagging strategies
- Right-sizing instances
- Spot instance strategies
- Cost allocation and budgets

MODULE 03

Git, GitHub & GitHub Actions

Week 3

3.1 Version Control with Git

- Introduction to code repositories & Git
- Clone, modify, and push code to remote repositories
- Code editors & productivity
- Branch best practices & scenarios for branching
- Commit messages: importance & effective writing
- Pulling code & pushing local changes
- Git tags: what and when
- Merging branches to master
- Git commit history
- Application release strategy using Git repos
- GitOps: introduction & implementation
- Trunk Based Development (TBD)
- Short-living vs Long-living branches

3.2 Collaboration with GitHub

- Introduction to GitHub
- Creating & optimizing GitHub profile for job marketing
- Creating repositories
- Pull requests & peer reviews

3.3 GitHub Actions (CI/CD)

- ✓ Overview & basic concepts: Workflows, Events, Jobs, Steps, Runners, Actions
- ✓ Setting up repository for GitHub Actions
- ✓ Writing workflow files & YAML syntax
- ✓ Triggering workflows (push, pull_request, etc.)
- ✓ Environment variables and secrets
- ✓ Job configuration: dependencies and parallel execution
- ✓ Using official and community actions
- ✓ Managing workflow dependencies and artifacts
- ✓ Conditional executions
- ✓ Deploying to Azure and Kubernetes
- ✓ Integrating with external services
- ✓ Notifications and best practices

MODULE 04

Ansible — Configuration Management

Week 4

4.1 Introduction

- ✓ Problems of shell scripting
- ✓ History of configuration management
- ✓ Ansible vs Chef vs Puppet

4.2 Ansible Basics

- ✓ Installation & version history
- ✓ Modules vs Collections
- ✓ Inventory management best practices
- ✓ ADHOC commands & SSH credentials

4.3 Ansible Playbooks

- ✓ XML vs JSON vs YAML
- ✓ Plays, multi-play playbooks, tasks
- ✓ DEBUG module

4.4 Variables

- ✓ Play-level, task-level, inventory, CLI variables
- ✓ Variable precedence

4.5 Roles & Advanced Concepts

- ✓ Role directory structure & best practices
- ✓ files, meta, templates directories
- ✓ Ansible Galaxy
- ✓ Conditions & loops
- ✓ Unified pull/push playbook

4.6 Ansible Ecosystem

- ✓ Ansible Tower/AWX
- ✓ GUI open source tools

MODULE 05

Terraform — Infrastructure as Code + Packer

Week 5

5.1 Introduction to IaC

- ✓ Advantages, alternatives, why Terraform

5.2 Terraform Language & Lifecycle

- ✓ HCL introduction, code structure, file extensions
- ✓ Installation
- ✓ Lifecycle: init, plan, apply, destroy

5.3 Variables & Functions

- ✓ Output block, inbuilt functions
- ✓ Variable block: tfvars, auto tfvars
- ✓ Variable precedence

5.4 Resources & Providers

- ✓ Resource block: attributes & arguments
- ✓ Provider block, data sources
- ✓ Resource reference

5.5 Loops & Inputs

- ✓ count, for_each
- ✓ Best input methods

5.6 State Management

- ✓ State file best practices
- ✓ Remote state backends

5.7 Modules & Conditions

- ✓ Local & vendor modules
- ✓ Module best practices
- ✓ Locals, conditions

5.8 Project Structure & Testing

- ✓ Versioning, multi-environment structure
- ✓ Best practices for real-time environments
- ✓ Terraform test cases

5.9 Packer for Azure VM Image Building NEW

- ✓ Introduction to Packer
- ✓ Packer templates for Azure VM Images
- ✓ Integration with Terraform

MODULE 06

Docker — Containerization

Week 6

6.1 Container Fundamentals

- ✓ Physical Server vs Virtual Server vs Container
- ✓ Container advantages & adoption
- ✓ Container management software types

6.2 Docker Introduction

- ✓ Docker architecture, ecosystem history
- ✓ Container runtimes
- ✓ Public Docker images

6.3 Docker Images & Dockerfiles

- ✓ Dockerfile instructions: FROM, RUN, COPY, ADD, CMD, ENTRYPOINT
- ✓ ADD vs COPY, CMD vs ENTRYPOINT
- ✓ Building images, image layers
- ✓ Multi-stage builds NEW
- ✓ Compare Docker image with Azure VM Image

6.4 Docker Registries

- ✓ Public vs private registries
- ✓ Pushing to registry
- ✓ Introduction to ACR
- ✓ Push Images to ACR
- ✓ Building images in pipelines

6.5 Running Containers

- ✓ Dynamic & static ports
- ✓ Volume mapping (single & multiple)
- ✓ Health checks
- ✓ Container inspection & management

MODULE 07

Kubernetes + Helm + Kustomize + ISTIO

Week 7–8

7.1 Kubernetes Fundamentals

- ✓ Orchestration introduction & importance
- ✓ Kubernetes vs Docker
- ✓ Pod advantages
- ✓ Architecture overview
- ✓ kubeconfig, Minikube setup
- ✓ Kubernetes YAML

7.2 Core Objects

- ✓ Pods: single/multi-container, labels, annotations, env vars
- ✓ ConfigMaps & Secrets
- ✓ ReplicaSet & Deployment
- ✓ Services: ClusterIP, LoadBalancer, NodePort

7.3 Kubernetes on Azure (AKS)

- ✓ AKS introduction & Terraform configuration
- ✓ OIDC in AKS, Node Pools, kubeconfig

7.4 Advanced Concepts

- ✓ Namespaces
- ✓ Azure Key Vault integration
- ✓ Service Accounts & Managed Identities
- ✓ Load Balancer exposure from AKS
- ✓ Nginx Ingress Controller
- ✓ Prometheus & Grafana on AKS

7.5 Karpenter / AKS Node Autoscaler — Node Autoscaling NEW

- ✓ Karpenter vs Cluster Autoscaler
- ✓ Provisioner configuration
- ✓ Node consolidation strategies

7.6 Resource Management NEW

- ✓ Requests & Limits
- ✓ Horizontal Pod Autoscaler (HPA)
- ✓ Vertical Pod Autoscaler (VPA)
- ✓ Pod Disruption Budgets

7.7 Kubernetes Gateway API NEW

- ✓ Gateway API vs Ingress
- ✓ HTTPRoute, GRPCRoute
- ✓ Traffic splitting & canary patterns

7.8 Helm Charts

- ✓ What is Helm? Architecture
- ✓ Chart structure: Chart.yaml, values.yaml, templates
- ✓ Creating charts from scratch
- ✓ Installing, upgrading, rolling back releases
- ✓ Helm repositories & best practices

7.9 Kustomize NEW

- ✓ Kustomize vs Helm
- ✓ Overlays for multi-environment management
- ✓ Integration with ArgoCD

7.10 Service Mesh — ISTIO & Cilium

- ✓ Why service mesh for microservices
- ✓ ISTIO sidecar injection (Envoy Proxy)
- ✓ TLS configuration, traffic management, observability
- ✓ Cilium — eBPF-based service mesh (industry trend) NEW
- ✓ ISTIO vs Cilium vs Linkerd comparison

MODULE 08

Monitoring & Observability

Week 9

8.1 Prometheus

- ✓ Monitoring & alerting importance
- ✓ Prometheus architecture: scrape, store, query
- ✓ Exporters: Node Exporter, cAdvisor
- ✓ Target configuration
- ✓ PromQL basics
- ✓ Alertmanager configuration
- ✓ Service discovery methods

8.2 Grafana

- ✓ Data visualization importance
- ✓ Configuring data sources
- ✓ Building dashboards & panels
- ✓ Visualization types
- ✓ Templating with variables
- ✓ Alerting setup

8.3 ELK Stack

- ✓ Centralized logging need
- ✓ Elasticsearch: index, document, shard, replica
- ✓ Logstash: inputs, filters, outputs
- ✓ Kibana: discover, visualizations, dashboards
- ✓ Use cases: troubleshooting, security analysis, performance

8.4 OpenTelemetry (OTEL) NEW

- ✓ Introduction to OpenTelemetry standard
- ✓ Traces, Metrics, and Logs unified
- ✓ OTEL Collector deployment
- ✓ Instrumenting applications
- ✓ Integration with Prometheus & Grafana

8.5 NewRelic

- ✓ APM introduction
- ✓ Key features: APM, infra monitoring, log management
- ✓ Agent integration
- ✓ Custom dashboards & alerts
- ✓ Troubleshooting with NewRelic

8.6 Fluent Bit / Fluentd NEW

- ✓ Modern log collection for Kubernetes
- ✓ Fluent Bit vs Logstash comparison
- ✓ Deploying as DaemonSet on AKS

8.7 SRE Practices NEW

- ✓ SLIs, SLOs, and Error Budgets
- ✓ Incident management & on-call
- ✓ Post-mortems & blameless culture
- ✓ Toil reduction strategies

MODULE 09

GitOps & CI/CD — ArgoCD

Week 10

9.1 ArgoCD

- ✓ What is GitOps? Architecture
- ✓ Installing ArgoCD in AKS
- ✓ Configuring ArgoCD
- ✓ Projects, YAML code, sync policies
- ✓ Best practices

9.2 Advanced CI/CD Patterns NEW

- ✓ Blue-green deployments
- ✓ Canary deployments
- ✓ Progressive delivery
- ✓ Feature flags overview

MODULE 10

Security & Quality

Week 11

10.1 SonarQube

- ✓ Continuous integration quality gates
- ✓ Installing SonarQube
- ✓ Sonar client setup, pipeline integration
- ✓ Bugs, vulnerabilities, code smells, coverage, duplication, debt
- ✓ Quality gates & profiles
- ✓ Pipeline failure on quality gate failure

10.2 HashiCorp Vault

- ✓ Secrets management use cases
- ✓ Vault CLI & setup
- ✓ KV Secrets, Transit Backend, Dynamic Secrets
- ✓ Authentication & audit logging
- ✓ Integration with projects, CI, CD, Kubernetes

10.3 Container & Infrastructure Security

- Trivy: container image scanning, K8s cluster scanning
- Lynis: Linux OS hardening & security auditing
- CIS benchmarks & CVE scanning

10.4 Policy as Code — OPA/Gatekeeper NEW

- Introduction to Open Policy Agent
- Gatekeeper for Kubernetes
- Writing Rego policies
- Enforcing policies in CI/CD and K8s

10.5 Software Supply Chain Security NEW

- SBOM (Software Bill of Materials) generation
- Image signing with Cosign/Sigstore
- Dependency vulnerability scanning
- SLSA framework overview

10.6 SAST/DAST Concepts NEW

- Static Application Security Testing
- Dynamic Application Security Testing
- CheckMarx SAST — static code analysis
- CheckMarx SCA — software composition analysis
- Integration in CI/CD pipelines

10.7 Service Mesh & Network Security NEW

- mTLS for service-to-service communication
- Web Firewall (Azure WAF) integration

MODULE 11

Python for DevOps

Week 12

11.1 Python Fundamentals

- History, versions (V2 vs V3), use cases
- Indentation, variables, data types
- Operators and expressions

11.2 Control Flow

- Conditions (if/elif/else)
- Loops (for, while), comprehensions

11.3 Functions

- Parameters, arguments, return values
- Scope and lifetime, doc strings, annotations
- Recursive functions

11.4 Data Structures

- Lists, Tuples, Dictionaries, Sets

11.5 Strings & Regular Expressions

- String manipulation, regular expressions

11.6 Modules for DevOps

- OS, Platform, Subprocess, Sys, Psutil modules
- Requests, Logging, Paramiko, Azure SDK for Python (azure-mgmt, azure-identity)
- Running Azure Functions with Python code

MODULE 12

AI for DevOps & AIOps ENTIRELY NEW

Week 13

12.1 AI-Assisted DevOps Workflows

- GitHub Copilot for Infrastructure as Code
- AI-assisted Terraform, Ansible, and K8s manifest generation
- Prompt engineering for DevOps automation
- AI-powered code review and testing

12.2 AIOps — Intelligent Operations

- Introduction to AIOps concepts
- AI-driven anomaly detection in monitoring
- Predictive alerting and auto-remediation
- Noise reduction in alert management
- ChatOps with AI assistants

12.3 LLM Infrastructure for DevOps

- Overview of deploying AI/ML models on Kubernetes
- GPU node management on AKS
- Model serving concepts (vLLM, TGI overview)
- Resource management for AI workloads

12.4 MLOps Pipeline Fundamentals

- CI/CD for ML models overview
- Model registries and versioning
- A/B testing for model deployments
- Monitoring ML model performance

MODULE 13

Capstone Project & Interview Prep

Week 14

13.1 End-to-End Project Implementation

- ✓ Complete microservices deployment on Azure
- ✓ Full CI/CD pipeline with security gates
- ✓ Monitoring and observability setup
- ✓ GitOps-based deployment

13.2 Real-World Troubleshooting Scenarios NEW

- ✓ Pod in CrashLoopBackOff — systematic debugging approach
- ✓ Terraform apply fails — state lock, drift, dependency issues
- ✓ CI/CD pipeline failure — build, test, deploy stage debugging
- ✓ Production CPU/Memory spike — identifying & resolving root cause
- ✓ 503 Service Unavailable — tracing from App Gateway to pod
- ✓ Docker image build failure — layer caching, multi-stage issues
- ✓ Kubernetes node NotReady — kubelet, resource, networking diagnosis
- ✓ Secrets leaked in Git — remediation & prevention with Vault
- ✓ ArgoCD out-of-sync — drift detection & resolution
- ✓ Prometheus alerts firing false positives — tuning thresholds & SLOs

13.3 Interview Preparation

- ✓ DevOps scenario-based questions
- ✓ Architecture design discussions
- ✓ Mock interviews with real-time feedback
- ✓ Resume review & ATS optimization

13.4 Certification Roadmap

- ✓ Microsoft Azure Administrator (AZ-104) — syllabus covers 80%+ topics
- ✓ Azure DevOps Engineer Expert (AZ-400) — syllabus covers 75%+ topics
- ✓ Certified Kubernetes Administrator (CKA) — syllabus covers 85%+ topics
- ✓ HashiCorp Terraform Associate — syllabus covers 90%+ topics
- ✓ Guidance on exam registration, study resources & preparation strategy