



Master in DevOps Engineering (MDE)

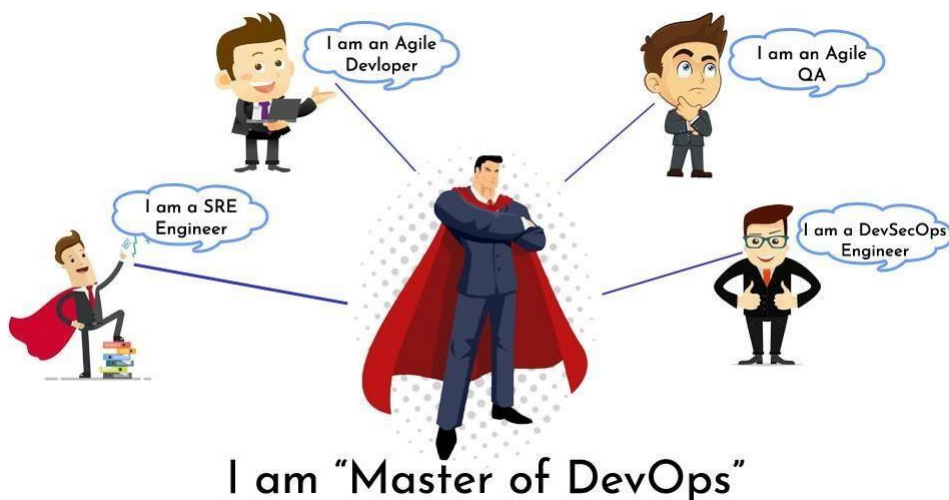
About DevOpsSchool

DevOpsSchool is a unit of "Cotocus PVT ltd" and a leading platform which helps IT organizations and professionals to learn all the emerging technologies and trend which helps them to learn and embrace all the skills, intelligence, innovation and transformation which requires to achieve the end result, quickly and efficiently. We provide over 40 specialized programs on DevOps, Cloud, Containers, Security, AI, ML and on Big data that are focused on industry requirement and each curriculum is developed and delivered by leading experts in each domain and aligned with the industry standards.

About Course

This "Masters in DevOps Engineering (MDE)" Program is the only course in the WORLD which can make you an expert and proficient Architect in DevOps, DevSecOps and Site Reliability Engineering (SRE) principles together. Our curriculum has been determined by comprehensive research on 10000+ job descriptions across the globe and epitome of 200+ years of industry experience.

"Master in DevOps Engineering" program is structured in a way, whether you are an experienced IT professional or a college graduate, this course will help you to integrate all the real-world experience with all the important tools, specialization and job-ready skills.





Co-ordinator – Akanksha

Call/WhatsApp: - +91 1800 889 7977

Mail Address: - contact@DevOpsSchool.com

Secondary contact – Patrick

Call/WhatsApp: - +91 7004 215 841

Mail Address:- contact@DevOpsSchool.com

Duration	120 Hrs
Mode	Online (Instructor-led, live & Interactive)
Projects (Real time scenario based)	3

FEATURES	DEVOPSSCHOOL	OTHERS
1 Course for All (DevOps/DevSecOps/SRE)	✓	✗
Faculty Profile Check	✓	✗
Lifetime Technical Support	✓	✗
Lifetime LMS access	✓	✗
Top 46 Tools	✓	✗
Interview KIT (Q&A)	✓	✗
Training Notes	✓	✗
Step by Step Web Based Tutorials	✓	✗
Training Slides	✓	✗
Training + Additional Videos	✓	✗



Training

Upon completion of this program you will get 360-degree understanding of DevOps, DevSecOps and SRE. This course will give you thorough learning experience in terms of understanding the concepts, mastering them thoroughly and applying them in real work environment.

Projects

You will be given industry level real time projects to work on and it will help you to differentiate yourself with multi-platform fluency, and have real-world experience with the most important tools and platforms.

Interview

As part of this, You would be given complete interview preparations kit, set to be ready for the DevOps hotseat. This kit has been crafted by 200+ years industry experience and the experiences of nearly 10000 DevOpsSchool's DevOps learners worldwide.



Agenda/Course outline of the MDE Training (DevOps/DevSecOps/SRE)

SDLC Models & Architecture with Agile, DevOps, SRE & DevSecOps, SOA & Micro services - Concept

- Let's Understand about Software Development Model
- Overview of Waterfall Development Model
- Challenges of Waterfall Development Model
- Overview of Agile Development Model
- Challenges of Agile Development Model
- Requirement of New Software Development Model
- Understanding an existing Pain and Waste in Current Software Development Model
- What is DevOps?
- Transition in Software development model
 - Waterfall -> Agile -> CI/CD -> DevOps -> DevSecOps
- Understand DevOps values and principles
- Culture and organizational considerations
- Communication and collaboration practices
- Improve your effectiveness and productivity
- DevOps Automation practices and technology considerations
- DevOps Adoption considerations in an enterprise environment
- Challenges, risks and critical success factors
- What is DevSecOps?
 - Let's Understand DevSecOps Practices and Toolsets.
- What is SRE?
 - Let's Understand SRE Practices and Toolsets.
- List of Tools to become Full Stack Developer/QA/SRE/DevOps/DevSecOps
- Microservices Fundamentals
- Microservices Patterns
 - Choreographing Services
 - Presentation components
 - Business Logic

- Database access logic
 - Application Integration
 - Modelling Microservices
 - Integrating multiple Microservices
 - Avoiding Breaking Changes
 - Choosing the right protocols
- Keeping it simple
 - Sync & Async
 - Dealing with legacy systems
 - Testing
 - What and When to test
 - Preparing for deployment
 - Monitoring Microservice Performance
 - Tools used for Microservices Demo using container



Platform - Operating Systems - Centos/Ubuntu & VirtualBox & Vagrant

- Installing CentOS7 and Ubuntu
- Accessing Servers with SSH
- Working at the Command Line
- Reading Files
- Using the vi Text Editor
- Piping and Redirection
- Archiving Files
- Accessing Command Line Help
- Understanding File Permissions
- Accessing the Root Account
- Using Screen and Script
- Overview of Hypervisor
- Introduction of VirtualBox
- Install VirtualBox and Creating CentOS7 and Ubuntu Vms
- Understanding Vagrant
- Basic Vagrant Workflow
- Advance Vagrant Workflow
- Working with Vagrant VMs
- The Vagrantfile
- Installing Nginx
- Provisioning
- Networking
- Sharing and Versioning Web Site Files
- Vagrant Share
- Vagrant Status
- Sharing and Versioning Nginx Config Files
- Configuring Synced Folders



Platform - Cloud - AWS - 4 hrs

- Introduction of AWS
- Understanding AWS infrastructure
- Understanding AWS Free Tier
- IAM: Understanding IAM Concepts
- IAM: A Walkthrough IAM
- IAM: Demo & Lab
- Computing:EC2: Understanding EC2 Concepts
- Computing:EC2: A Walkthrough EC2
- Computing:EC2: Demo & Lab
- Storage:EBS: Understanding EBS Concepts
- Storage:EBS: A Walkthrough EBS
- Storage:EBS: Demo & Lab
- Storage:S3: Understanding S3 Concepts
- Storage:S3: A Walkthrough S3
- Storage:S3: Demo & Lab
- Storage:EFS: Understanding EFS Concepts
- Storage:EFS: A Walkthrough EFS
- Storage:EFS: Demo & Lab
- Database:RDS: Understanding RDS MySql Concepts
- Database:RDS: A Walkthrough RDS MySql
- Database:RDS: Demo & Lab
- ELB: Elastic Load Balancer Concepts
- ELB: Elastic Load Balancer Implementation
- ELB: Elastic Load Balancer: Demo & Lab
- Networking:VPC: Understanding VPC Concepts
- Networking:VPC: Understanding VPC components
- Networking:VPC: Demo & Lab

Platform - Containers - Docker - 4 hrs

- What is Containerization?
- Why Containerization?
- How Docker is good fit for Containerization?
- How Docker works?
- Docker Architecture
- Docker Installations & Configurations
- Docker Components
- Docker Engine
- Docker Image
- Docker Containers
- Docker Registry
- Docker Basic Workflow
- Managing Docker Containers
- Creating our First Image
- Understanding Docker Images
- Creating Images using Dockerfile
- Managing Docker Images
- Using Docker Hub registry
- Docker Networking
- Docker Volumes
- Deepdive into Docker Images
- Deepdive into Dockerfile
- Deepdive into Docker Containers
- Deepdive into Docker Networks
- Deepdive into Docker Volumes
- Deepdive into Docker Volume
- Deepdive into Docker CPU and RAM allocations
- Deepdive into Docker Config
- Docker Compose Overview
- Install & Configure Compose
- Understanding Docker Compose Workflow
- Understanding Docker Compose Services
- Writing Docker Compose Yaml file
- Using Docker Compose Commands
- Docker Compose with Java Stake
- Docker Compose with Rails Stake
- Docker Compose with PHP Stake
- Docker Compose with Nodejs Stake



Planning and Designing - Jira & Confluence (2 + 2 = 4 hrs)

Jira

- Overview of Jira
- Use cases of Jira
- Architecture of Jira
- Installation and Configuration of Jira in Linux
- Installation and Configuration of Jira in Windows
- Jira Terminologies
- Understanding Types of Jira Projects
- Working with Projects
- Working with Jira Issues
- Adding Project Components and Versions
- Use Subtasks to Better Manage and Structure Your Issues
- Link Issues to Other Resources
- Working in an Agile project
- Working with Issues Types by Adding/Editing/Deleting
- Working with Custom Fields by Adding/Editing/Deleting
- Working with Screens by Adding/Editing/Deleting
- Searching and Filtering Issues
- Working with Workflow basic
- Introduction of Jira Plugins and Addons.
- Jira Integration with Github

Confluence

- Exploring Confluence benefits and resources
- Configuring Confluence
- Navigating the dashboard, spaces, and pages
- Creating users and groups
- Creating pages from templates and blueprints
- Importing, updating, and removing content
- Giving content feedback
- Watching pages, spaces, and blogs
- Managing tasks and notifications
- Backing up and restoring a site
- Admin tasks
 - Add/Edit/Delete new users
 - Adding group and setting permissions
 - Managing user permissions
 - Managing addons or plugins
 - Customizing confluence site

- Installing Confluence
 - Evaluation options for Confluence
 - Supported platforms
 - Installing Confluence on Windows
 - Activating Confluence trial license
 - Finalizing Confluence Installation

Backend Programming Language 2 - Python/Flask with MySQL DB - 4 hrs.

- Planning - Discuss some of the Small Project Requirement which include Login/Registration with Some Students records CRUD operations.
- Design a Method --> Classes -> Interface using Core Python
 - Fundamental of Core Python with Hello-world Program with Method -->Classes
- Coding in Flask using HTML - CSS - JS - MySQL
 - Fundamental of Flask Tutorial of Hello-World App
- UT - 2 Sample unit Testing using Python test
- Package a Python App
- AT - 2 Sample unit testing using Selenium

Deploy to Some env.

.....END OF THE CLASS

Technology Demonstration

- Software Planning and Designing using JAVA
- Core Python
- Flask
- mySql
- pytest
- Selenium
- HTML
- CSS
- Js.



Source Code Versioning - Git using GitHub - 4 hrs.

- Introduction of Git
- Installing Git
- Configuring Git
- Git Concepts and Architecture
- How Git works?
- The Git workflow
- Working with Files in Git
 - Adding files
 - Editing files
 - Viewing changes with diff
 - Viewing only staged changes
 - Deleting files
 - Moving and renaming files
 - Making Changes to Files
- Undoing Changes
 - Reset
 - Revert
- Amending commits
- Ignoring Files
- Branching and Merging using Git
- Working with Conflict Resolution
- Comparing commits, branches and workspace
- Working with Remote Git repo using Github
- Push - Pull - Fetch using Github
- Tagging with Git

Code Analysis & Securing Code (SAST) - SonarQube & Coverity Scan & Snyk

- What is SonarQube?
- Benefits of SonarQube?
- Alternative of SonarQube
- Understanding Various License of SonarQube
- Architecture of SonarQube
- How SonarQube works?
- Components of SonarQube
- SonarQube runtime requirements
- Installing and configuring SonarQube in Linux
- Basic Workflow in SonarQube using Command line
- Working with Issues in SonarQube
- Working with Rules in SonarQube
- Working with Quality Profiles in SonarQube
- Working with Quality Gates in SonarQube
- Deep Dive into SonarQube Dashboard
- Understanding Seven Axis of SonarQube Quality
- Workflow in SonarQube with Maven Project
- Workflow in SonarQube with Gradle Project
- OWASP Top 10 with SonarQube

Build Management - Maven and Gradle - 2 + 2 = 4 hrs

Maven

- Introduction to Apache Maven
- Advantage of Apache Maven over other build tools
- Understanding the Maven Lifecycle and Phase
- Understanding the Maven Goals
- Understanding the Maven Plugins
- Understanding the Maven Repository
- Understanding and Maven Release and Version
- Prerequisite and Installing Apache Maven
- Understanding and using Maven Archetypes
- Understanding Pom.xml and Setting.xml
- Playing with multiples Maven Goals
- Introducing Maven Dependencies
- Introducing Maven Properties
- Introducing Maven Modules
- Introducing Maven Profile
- Introducing Maven Plugins
- How can Maven benefit my development process?
- How do I setup Maven?
- How do I make my first Maven project?
- How do I compile my application sources?
- How do I compile my test sources and run my unit tests?
- How do I create a JAR and install it in my local repository?
- How do I use plugins?
- How do I add resources to my JAR?
- How do I filter resource files?
- How do I use external dependencies?
- How do I deploy my jar in my remote repository?
- How do I create documentation?
- How do I build other types of projects?
- How do I build more than one project at once?

Gradle

- What is Gradle?
- Why Gradle?
- Installing and Configuring Gradle
- Build Java Project with Gradle
- Build C++ Project with Gradle
- Build Python Project with Gradle
- Dependency Management in Gradle
- Project Structure in Gradle
- Gradle Tasks
- Gradle Profile and Cloud
- Gradle Properties
- Gradle Plugins

Package Management - Packer & Artifactory - 2 + 2 = 4hrs

Artifactory

- Artifactory Overview
- Understanding a role of Artifactory in DevOps
- System Requirements
- Installing Artifactory in Linux
- Using Artifactory
- Getting Started
- General Information
- Artifactory Terminology
- Artifactory Repository Types
- Artifactory Authentication
- Deploying Artifacts using Maven
- Download Artifacts using Maven
- Browsing Artifactory
- Viewing Packages
- Searching for Artifacts
- Manipulating Artifacts

Packer

- Getting to Know Packer
 - What is Packer?
 - Save what is Packer?

- Installing Packer
- Save Installing Packer

- The Packer workflow and components
- Save The Packer workflow and components
- The Packer CLI
- Save The Packer CLI

- Baking a Website Image for EC2
 - Select an AWS AMI base
 - Save Select an AWS AMI base
 - Automate AWS AMI base build
 - Save Automate AWS AMI base build
 - Using build variables
 - Save Using build variables
 - Provision Hello World
 - Save Provision Hello World
 - Provision a basic site
 - Save Provision a basic site

- Customization with a Config Management Tool
 - Simplify provisioning with a config tool
 - Save Simplify provisioning with a config tool
 - Use ansible to install the webserver
 - Save Use ansible to install the webserver
 - Debugging
 - Save Debugging

- Building Hardened Images
 - Use Ansible modules to harden our image
 - Save Use Ansible modules to harden our image
 - Baking a Jenkins image
 - Save Baking a Jenkins image

- Building a Pipeline for Packer Image
 - Validate Packer templates
 - Save Validate Packer templates
 - Create a manifest profile
 - Save Create a manifest profile
 - Testing
 - Save Testing
 - CI pipeline
 - Save CI pipeline

Unit Testing & Acceptance Testing & Coverage - Junit & Selenium & Jmeter & Jacoco - 4hrs

- Junit Fundamental - 1 Hour
 - What is Unit Testing
 - Tools for Unit Testing
 - What is Junit?
 - How to configure Junit?
 - Writing Basic Junit Test cases
 - Running Basic Junit Test cases
 - Junit Test Results

- Selenium Fundamental -2 Hours
 - Introduction To Selenium
 - Components of Selenium
 - Selenium IDE
 - Selenium Webdriver
 - Selenium Grid
 - Installing and Configuring Selenium
 - Working with Selenium IDE
 - Working With Selenium Webdriver with Java Test Case
 - Setup and Working with Selenium Grid

- Jacoco - 1 Hours
 - Overview of Code Coverage Process
 - Introduction of Jacoco
 - How Jacoco works!
 - How to install Jaoco?
 - Setup testing Environment with Jacoco
 - Create test data files using Jacoco and Maven
 - Create a Report using Jacoco
 - Demo - Complete workflow of Jacoco with Maven and Java Project

Configuration & Deployment Management - Ansible - 4 hrs

- Overflow of Configuration Management
- Introduction of Ansible
- Ansible Architecture
- Let's get started with Ansible
- Ansible Authentication & Authorization
- Let's start with Ansible Adhoc commands
- Let's write Ansible Inventory
- Let's write Ansible Playbook
- Working with Popular Modules in Ansible
- Deep Dive into Ansible Playbooks
- Working with Ansible Variables
- Working with Ansible Template
- Working with Ansible Handlers
- Roles in Ansible
- Ansible Galaxy

Container Orchestration - Kubernetes & Helm Introduction - 4 hrs

Kubernetes

- Understanding the Need of Kubernetes
- Understanding Kubernetes Architecture
- Understanding Kubernetes Concepts
- Kubernetes and Microservices
- Understanding Kubernetes Masters and its Component
 - kube-apiserver
 - etcd
 - kube-scheduler
 - kube-controller-manager
- Understanding Kubernetes Nodes and its Component
 - kubelet
 - kube-proxy
 - Container Runtime
- Understanding Kubernetes Addons

- DNS
- Web UI (Dashboard)
- Container Resource Monitoring
- Cluster-level Logging
- Understand Kubernetes Terminology
- Kubernetes Pod Overview
- Kubernetes ReplicationController Overview
- Kubernetes Deployment Overview
- Kubernetes Service Overview
- Understanding Kubernetes running environment options
- Working with first Pods
- Working with first ReplicationController
- Working with first Deployment
- Working with first Services
- Introducing Helm
- Basic working with Helm

Infrastructure Coding - Terraform - 4 hrs

- Deploying Your First Terraform Configuration
 - Introduction
 - What's the Scenario?
 - Terraform Components
- Updating Your Configuration with More Resources
 - Introduction
 - Terraform State and Update
 - What's the Scenario?
 - Data Type and Security Groups
- Configuring Resources After Creation
 - Introduction
 - What's the Scenario?
 - Terraform Provisioners
 - Terraform Syntax
- Adding a New Provider to Your Configuration

- Introduction
- What's the Scenario?
- Terraform Providers
- Terraform Functions
- Intro and Variable
- Resource Creation
- Deployment and Terraform Console
- Updated Deployment and Terraform Commands

Continuous Integration - Jenkins - 4 hrs.

- Let's understand Continuous Integration
 - What is Continuous Integration
 - Benefits of Continuous Integration
 - What is Continuous Delivery
 - What is Continuous Deployment
 - Continuous Integration Tools
- What is Jenkins
 - History of Jenkins
 - Jenkins Architecture
 - Jenkins Vs Jenkins Enterprise
 - Jenkins Installation and Configurations
- Jenkins Dashboard Tour
 - Understand Freestyle Project
 - Freestyle General Tab
 - Freestyle Source Code Management Tab
 - Freestyle Build Triggers Tab
 - Freestyle Build Environment
 - Freestyle Build
 - Freestyle Post-build Actions
 - Manage Jenkins
 - My Views & Credentials
 - People & Build History
- Creating a Simple Job
 - Simple Java and Maven Based Application

- Simple Java and Gradle Based Application
 - Simple DOTNET and MSBuild Based Application
- Jobs Scheduling in Jenkins
 - Manually Building
 - Build Trigger based on fixed schedule
 - Build Trigger by script
 - Build Trigger Based on pushed to git
- Useful Jobs Configuration
 - Jenkins Jobs parameterised
 - Execute concurrent builds
 - Jobs Executors
 - Build Other Projects
 - Build after other projects are built
 - Throttle Builds
- Jenkins Plugins
 - Installing a Plugin
 - Plugin Configuration
 - Updating a Plugin
 - Plugin Wiki
 - Top 20 Useful Jenkins Plugins
 - Using Jenkins Plugins Best Practices
- Jenkins Node Management
 - Adding a Linux Node
 - Adding a Windows Nodes
 - Nodes Management using Jenkins & Jenkins Nodes High Availability
- Jenkins Integration with other tools
 - Jira
 - Git
 - SonarQube
 - Maven
 - Junit
 - Ansible
 - Docker

 - AWS

- Jacoco
- Coverity
- Selenium
- Gradle
- Reports in Jenkins
 - Junit Report
 - SonarQube Reports
 - Jacoco Reports
 - Coverity Reports
 - Selenium Reports
 - Test Results
 - Cucumber Reports
- Jenkins Node Management
 - Adding a Linux Node
 - Adding a Windows Nodes
 - Nodes Management using Jenkins
 - Jenkins Nodes High Availability
- Notification & Feedback in Jenkins
 - CI Build Pipeline & Dashboard
 - Email Notification
 - Advance Email Notification
 - Slack Notification
- Jenkins Advance – Administrator
 - Security in Jenkins
 - Authorization in Jenkins
 - Authentication in Jenkins
 - Managing folder/subfolder
 - Jenkins Upgrade
 - Jenkins Backup
 - Jenkins Restore
 - Jenkins Command Line



Infrastructure Monitoring Tool 2 - Datadog - 4 hrs.

Datadog – 3 hrs.

- Introduction
 - Introduction to Datadog
 - Datadog installation
 - Grafana with Datadog Installation
- Monitoring
 - Introduction to Monitoring
 - Client Libraries
 - Pushing Metrics
 - Querying
 - Service Discovery
 - Exporters
- Alerting
 - Introduction to Alerting
 - Setting up Alerts
- Internals
 - Datadog Storage
 - Datadog Security
 - TLS & Authentication on Datadog Server
 - Mutual TLS for Datadog Targets
- Use Cases
 - Monitoring a web application
 - Calculating Apdex score
 - Cloudwatch Exporter
 - Grafana Provisioning
 - Consul Integration with Datadog
 - EC2 Auto Discovery
 - Installing on Mac
 - Installing using Docker
 - Building from source
 - Upgrading

- Administration
 - Configuration
 - Authentication
 - Permissions
 - Grafana CLI
 - Internal metrics
 - Provisioning & Troubleshooting

Log Monitoring Tool 1 - Splunk - 4 hrs

- What Is Splunk?
 - Overview
 - Machine Data
 - Splunk Architecture
 - Careers in Splunk
- Setting up the Splunk Environment
 - Overview
 - Splunk Licensing
 - Getting Splunk
 - Installing Splunk
 - Adding Data to Splunk
- Basic Searching Techniques
 - Adding More Data
 - Search in Splunk
 - Demo: Splunk Search
 - Splunk Search Commands
 - Splunk Processing Language
 - Splunk Reports
 - Reporting in Splunk
 - Splunk Alerts
 - Alerts in Splunk
- Enterprise Splunk Architecture
 - Overview
 - Forwarders

- Enterprise Splunk Architecture
- Installing Forwarders
- Installing Forwarders
- Troubleshooting Forwarder Installation
- Splunking for DevOps and Security
 - Splunk in DevOps
 - DevOps Demo
 - Splunk in Security & Enterprise Use Cases
- Application Development in Splunkbase
 - What Is Splunkbase?
 - Navigating the Splunkbase
 - Creating Apps for Splunk
 - Benefits of Building in Splunkbase
- Splunking on Hadoop with Hunk
 - What Is Hadoop?
 - Running HDFS Commands
 - What Is Hunk?
 - Installing Hunk
 - Moving Data from HDFS to Hunk
- Composing Advanced Searches
 - Splunk Searching
 - Introduction to Advanced Searching
 - Eval and Fillnull Commands
 - Other Splunk Command Usage
 - Filter Those Results! & The Search Job Inspector
- Creating Search Macros
 - What Are Search Macros?
 - Using Search Macros within Splunk
 - Macro Command Options and Arguments
 - Other Advanced Searching within Splunk

Performance & RUM Monitoring - NewRelic - 4 hrs

- Introduction and Overview of NewRelic
 - What is Application Performance Management?
 - Understanding a need of APM
 - Understanding transaction traces
 - What is Application Performance?
 - APM Benefits
 - APM Selection Criteria
 - Why NewRelic is best for APM?
 - What is NewRelic APM? & How does NewRelic APM work?
 - NewRelic Architecture & Terminology
- Installing and Configuring NewRelic APM Agents for Application
 - Register a NewRelic Trial account
 - Installing a JAVA Agent to Monitor your Java Application
 - Installing a PHP Agent to Monitor your PHP Application
 - Installing New Relic Agent for .NET Framework Application
 - Installing a Docker based Agent to Monitor your Docker based Application
 - Understanding of NewRelic Configuration settings of newrelic.yml
 - Understanding of NewRelic Agent Configuration settings
- Working with NewRelic Dashboard
 - Understanding a transactions
 - Understanding Apdex and Calculating and Setting Apdex Threshold
 - Understanding Circuit break
 - Understanding Throughput
 - Newrelic default graphs
 - Understanding and Configuring Service Maps
 - Understanding and Configuring JVM
 - Understanding Error Analytics
 - Understanding Violations
 - Understanding and Configuring Deployments
 - Understanding and Configuring Thread Profiler
 - Depp Dive into Transaction Traces

- Profiling with New Relic
- Creating and managing Alerts
- Working with Incidents
- Sending NewRelic Alerts to Slack
- Assessing the quality of application deployments
- Monitoring using Newrelic
 - View your applications index
 - APM Overview page
 - New Relic APM data in Infrastructure
 - Transactions page
 - Databases and slow queries
 - Viewing slow query details & External services page
 - Agent-specific UI & Viewing the transaction map
- Deep Dive into Newrelic Advance
 - Newrelic transaction alerts
 - Configure abnd Troubleshoot and Cross Application Traces
 - NewRelic Service Level Agreements
 - Troubleshooting NewRelic
 - Understanding and Configuring NewRelic X-Ray Sessions
 - Deep Dive into NewRelic Agent Configuration
 - Adding Custom Data with the APM Agent
 - Extending Newrelic using Plugins
 - Finding and Fixing Application Performance Issues with New Relic APM
 - Setting up database monitoring using Newrelic APM
 - Setting up and Configuring Newrelic Alerts
- Working with NewRelic Performance Reports
 - Availability report
 - Background jobs analysis report
 - Capacity analysis report
 - Database analysis report
 - Host usage report
 - Scalability analysis report
 - Web transactions analysis report
 - Weekly performance report

Webserver - Apache HTTP & Nginx

Apache HTTP

- Introduction to web server
- Install Apache on CentOS 7.4
- Enable Apache to automatically start when system boot
- Configure the firewall service
- Where is Apache?
- Directory structure
 - Apache directory structure
 - Configuration file
 - Create your first page
- Virtual hosts
 - Setting up the virtual host - name based
 - Setting up the virtual host - port based
- Using aliases and redirecting
- Configuring an alias for a url
- Redirects
- Logging
 - The error log
 - The access log
 - Custom log
 - Log rotation
- Security
 - Basic Security - Part 1
 - Basic Security - Part 2
 - Set up TLS/SSL for free
 - Basic authentication
 - Digest authentication
 - Access Control
 - .htaccess (Administrator Side)
 - .htaccess (User Side)
 - Install and Configure antivirus
 - Mitigate dos attacks - mod_evasive
- Apache Performance and Troubleshooting
 - Apache Multi-Processing Modules (MPMs)
 - Adjusting httpd.conf - Part 1
 - Adjusting httpd.conf - Part 2
 - Troubleshoot Apache (Analyze Access Log) - Part 1
 - Troubleshoot Apache (Analyze Access Log) - Part 2
 - Use Apachetop to monitor web server traffic

Nginx

- Overview
 - Introduction
 - About NGINX
 - NGINX vs Apache
 - Test your knowledge
- Installation
 - Server Overview
 - Installing with a Package Manager
 - Building Nginx from Source & Adding Modules
 - Adding an NGINX Service
 - Nginx for Windows
 - Test your knowledge
- Configuration
 - Understanding Configuration Terms
 - Creating a Virtual Host
 - Location blocks
 - Variables
 - Rewrites & Redirects
 - Try Files & Named Locations
 - Logging
 - Inheritance & Directive types
 - PHP Processing
 - Worker Processes
 - Buffers & Timeouts
 - Adding Dynamic Modules
 - Test your knowledge
- Performance
 - Headers & Expires
 - Compressed Responses with gzip
 - FastCGI Cache
 - HTTP2
 - Server Push
- Security
 - HTTPS (SSL)
 - Rate Limiting
 - Basic Auth
 - Hardening Nginx
 - Test your knowledge
 - Let's Encrypt - SSL Certificates

Multi-cluster Kubernetes orchestration platform - Rancher

- Multi-cluster management
 - Rancher provides a unified interface for managing multiple Kubernetes clusters across different environments, including on-premises, cloud, and hybrid.
- Centralized administration
 - With Rancher, you can manage user access, security policies, and cluster settings from a central location, making it easier to maintain a consistent and secure deployment across all clusters.
- Automated deployment
 - Rancher streamlines the application deployment process by providing built-in automation tools that allow you to deploy applications to multiple clusters with just a few clicks.
- Monitoring and logging
 - Rancher provides a built-in monitoring and logging system that enables you to monitor the health and performance of your applications and clusters in real-time.
- Application catalog
 - Rancher offers a curated catalog of pre-configured application templates that enable you to deploy and manage popular applications such as databases, web servers, and messaging queues.
- Scalability and resilience
 - Rancher is designed to be highly scalable and resilient, enabling you to easily add new clusters or nodes to your deployment as your needs grow.
- Extensibility
 - Rancher provides an open API and a rich ecosystem of plugins and extensions, enabling you to customize and extend the platform to meet your specific needs.

Services mesh Data planes & Control Planes - Envoy & Istio

Envoy

Data Plane

- Envoy is a high-performance proxy that is deployed as a sidecar to each microservice in the infrastructure.
- Envoy manages all inbound and outbound traffic for the microservice and provides features like load balancing, circuit breaking, and health checks.
- Envoy can also be used as a standalone proxy outside of a service mesh architecture.

Control Plane:

- Envoy does not have a built-in control plane.

- It can be integrated with other service mesh management solutions like Istio, Consul, or Linkerd, which provide a central point of management for the Envoy proxies.
- These control planes enable features like traffic management, security, and observability.

Istio:

Data Plane:

- Istio uses Envoy as its data plane, which means that each microservice has an Envoy sidecar proxy that manages the inbound and outbound traffic for that service.
- Envoy is configured and managed by Istio's control plane components.

Control Plane:

- Istio provides a built-in control plane that includes the following components:
- Pilot: responsible for managing the configuration of the Envoy proxies and enabling features like traffic routing and load balancing.
- Mixer: provides policy enforcement, telemetry collection, and access control for the microservices in the service mesh.
- Citadel: responsible for managing the security of the service mesh, including mutual TLS encryption and identity-based access control.

Network configurations and Service Discovery - Consul

Network Configurations:

- Consul provides a central service registry that keeps track of all the services in the infrastructure.
- Each microservice in the infrastructure registers itself with Consul, providing information like its IP address, port, and health status.
- Consul also supports multiple datacenters, allowing for the deployment of services across different regions or availability zones.
- Consul provides a DNS interface that can be used to discover services in the infrastructure. Applications can use this interface to resolve service names to IP addresses and connect to the appropriate service.

Service Discovery:

- Consul provides a service discovery mechanism that enables microservices to discover and communicate with each other.
- Consul supports different service discovery methods, including DNS, HTTP, and gRPC.
- Consul can perform health checks on the services in the infrastructure to ensure that they are functioning properly. If a service fails a health check, it is removed from the service registry until it is healthy again.

- Consul also supports service segmentation, allowing services to be grouped into logical subsets based on tags or other attributes. This enables more fine-grained control over service discovery and traffic routing.

Securing Credentials - HashiCorp Vault

Secret Storage:

- Vault provides a secure storage mechanism for sensitive data, including credentials, API keys, and other secrets.
- Vault uses encryption and access control policies to ensure that secrets are protected both at rest and in transit.
- Vault supports different storage backends, including disk, cloud storage, and key management systems.

Authentication:

- Vault provides several authentication methods that can be used to validate user or machine identity.
- These methods include LDAP, Active Directory, Kubernetes, and token-based authentication.
- Vault also supports multi-factor authentication (MFA) to provide an additional layer of security.

Access Control:

- Vault provides fine-grained access control policies that can be used to restrict access to specific secrets or resources.
- These policies can be based on user or machine identity, time of day, and other factors.
- Vault supports role-based access control (RBAC) and attribute-based access control (ABAC) policies.

Encryption:

- Vault provides end-to-end encryption for all secrets stored in its storage backend.
- Vault uses encryption keys that are stored separately from the secrets themselves, providing an additional layer of security.
- Vault supports different encryption algorithms and key management systems.

Auditing and Logging:

- Vault provides detailed auditing and logging capabilities that can be used to track access to secrets and detect potential security threats.

- Vault logs all user and system activity, including authentication events, secret access, and configuration changes.
- Vault also supports integration with popular logging and monitoring tools.

Infrastructure Monitoring Tool - Prometheus with Grafana

Prometheus

- Introduction
- Introduction to Prometheus
- Prometheus installation
- Grafana with Prometheus Installation
- Monitoring
- Introduction to Monitoring
- Client Libraries
- Pushing Metrics
- Querying
- Service Discovery
- Exporters
- Alerting
- Introduction to Alerting
- Setting up Alerts
- Internals
- Prometheus Storage
- Prometheus Security
- TLS & Authentication on Prometheus Server
- Mutual TLS for Prometheus Targets
- Use Cases
- Monitoring a web application
- Calculating Apdex score
- Cloudwatch Exporter
- Grafana Provisioning
- Consul Integration with Prometheus
- EC2 Auto Discovery

Grafana

- Installation
- Installing on Ubuntu / Debian
- Installing on Centos / Redhat
- Installing on Windows
- Installing on Mac
- Installing using Docker
- Building from source
- Upgrading
- Administration
- Configuration
- Authentication
- Permissions
- Grafana CLI
- Internal metrics
- Provisioning
- Troubleshooting

Log Monitoring Tool - Elasticsearch Logstash Kibana(ELK stake)

- Introduction to Elasticsearch
- Overview of the Elastic Stack (ELK+)
- Elastic Stack
- Architecture of Elasticsearch
- Nodes & Clusters
- Indices & Documents
- A word on types
- Another word on types
- Sharding
- Replication
- Keeping replicas synchronized
- Searching for data

- Distributing documents across shards
- Installing Elasticsearch & Kibana
- Running Elasticsearch & Kibana in Elastic Cloud
- Installing Elasticsearch on Mac/Linux
- Using the MSI installer on Windows
- Installing Elasticsearch on Windows
- Configuring Elasticsearch
- Installing Kibana on Mac/Linux
- Installing Kibana on Windows
- Configuring Kibana
- Kibana now requires data to be available
- Introduction to Kibana and dev tools
- Managing Documents
- Creating an index
- Adding documents
- Retrieving documents by ID
- Replacing documents
- Updating documents
- Scripted updates
- Upserts
- Deleting documents
- Deleting indices
- Batch processing
- Importing test data with cURL
- Exploring the cluster
- Mapping
- Introduction to mapping
- Dynamic mapping
- Meta fields
- Field data types
- Adding mappings to existing indices
- Changing existing mappings
- Mapping parameters

- Adding multi-fields mappings
- Defining custom date formats
- Picking up new fields without dynamic mapping
- Analysis & Analyzers
- Introduction to the analysis process
- A closer look at analyzers
- Using the Analyze API
- Understanding the inverted index
- Analyzers
- Overview of character filters
- Overview of tokenizers
- Overview of token filters
- Overview of built-in analyzers
- Configuring built-in analyzers and token filters
- Creating custom analyzers
- Using analyzers in mappings
- Adding analyzers to existing indices
- A word on stop words
- Introduction to Searching
- Search methods
- Searching with the request URI
- Introducing the Query DSL
- Understanding query results
- Understanding relevance scores
- Debugging unexpected search results
- Query contexts
- Full text queries vs term level queries
- Basics of searching
- Term Level Queries
- Introduction to term level queries
- Searching for a term
- Searching for multiple terms

- Retrieving documents based on IDs
- Matching documents with range values
- Working with relative dates (date math)
- Matching documents with non-null values
- Matching based on prefixes
- Searching with wildcards
- Searching with regular expressions
- Term Level Queries
- Full Text Queries
- Introduction to full text queries
- Flexible matching with the match query
- Matching phrases
- Searching multiple fields
- Full Text Queries
- Adding Boolean Logic to Queries
- Introduction to compound queries
- Querying with boolean logic
- Debugging bool queries with named queries
- How the "match" query works

Incident Response Tool - PagerDuty & Opsgenie

Alert Management:

- Both PagerDuty and Opsgenie provide powerful alert management capabilities, allowing teams to configure alerts based on specific criteria, such as event severity, priority, and more.
- Alerts can be sent to multiple channels, including email, SMS, voice, and mobile push notifications.
- Both tools also provide support for escalation policies, allowing teams to ensure that critical alerts are addressed promptly.

Incident Management:

- Both PagerDuty and Opsgenie provide incident management capabilities, allowing teams to track incidents and collaborate on resolving them.

- Incident management features include creating incidents, adding notes, assigning owners, and tracking status changes.
- Both tools also provide support for incident timelines, allowing teams to visualize the progress of an incident over time.

Integration:

- Both PagerDuty and Opsgenie provide extensive integration capabilities, allowing teams to integrate with a wide range of tools and technologies.
- Integrations include popular monitoring tools, such as Nagios, New Relic, and AWS CloudWatch, as well as IT service management (ITSM) tools like JIRA and ServiceNow.
- Both tools also provide REST APIs for custom integrations.

Analytics and Reporting:

- Both PagerDuty and Opsgenie provide analytics and reporting capabilities, allowing teams to track performance metrics and identify areas for improvement.
- Analytics and reporting features include incident duration, resolution times, and other key performance indicators (KPIs).
- Both tools also provide support for custom dashboards and reports.

Automation:

- Both PagerDuty and Opsgenie provide automation capabilities, allowing teams to automate repetitive tasks and streamline incident response processes.
- Automation features include auto-acknowledgment of alerts, auto-escalation of incidents, and auto-remediation of issues.
- Both tools also provide support for scripting and custom automation workflows.

Production Env Job scheduler and Run Book Automation - RunDeck

Job Scheduling:

- RunDeck provides powerful job scheduling capabilities, allowing teams to schedule jobs based on specific criteria, such as time, date, and recurrence.
- Jobs can be executed on multiple platforms, including Windows, Linux, and macOS.
- RunDeck also provides support for job dependencies, allowing teams to ensure that jobs are executed in the correct order.

Run Book Automation:

- RunDeck provides run book automation capabilities, allowing teams to automate repetitive tasks and streamline operations.
- Run book automation features include executing commands, scripts, and workflows on multiple systems, as well as orchestrating complex processes across multiple systems.
- RunDeck also provides support for auditing and logging, allowing teams to track changes and monitor system activity.

Integration:

- unDeck provides extensive integration capabilities, allowing teams to integrate with a wide range of tools and technologies.
- Integrations include popular configuration management tools, such as Ansible and Puppet, as well as monitoring tools like Nagios and Zabbix.
- RunDeck also provides REST APIs for custom integrations.

Access Control:

- RunDeck provides access control capabilities, allowing teams to control who can access and execute jobs and workflows.
- Access control features include role-based access control (RBAC), LDAP integration, and multi-factor authentication (MFA).
- RunDeck also provides support for audit logging, allowing teams to track user activity and changes to system configurations.

Notifications and Reporting:

- RunDeck provides notifications and reporting capabilities, allowing teams to track performance metrics and identify areas for improvement.
- Notifications and reporting features include job execution status, error notifications, and custom reports.
- RunDeck also provides support for custom dashboards and reports.

Application Performance Monitoring - Appdynamics

Application Performance Monitoring:

- AppDynamics provides powerful application performance monitoring capabilities, allowing teams to monitor the performance of their applications in real-time.
- APM features include application topology maps, transaction tracing, code-level diagnostics, and performance baselines.

- AppDynamics also provides support for identifying and troubleshooting performance issues, such as slow database queries, inefficient code, and memory leaks.

End-User Monitoring:

- AppDynamics provides end-user monitoring capabilities, allowing teams to track the performance of their applications from the end-user perspective.
- End-user monitoring features include real-user monitoring, synthetic monitoring, and business transaction monitoring.
- AppDynamics also provides support for identifying and troubleshooting end-user issues, such as slow page load times and errors.

Infrastructure Monitoring:

- AppDynamics provides infrastructure monitoring capabilities, allowing teams to monitor the health and performance of their infrastructure.
- Infrastructure monitoring features include server monitoring, container monitoring, and cloud monitoring.
- AppDynamics also provides support for identifying and troubleshooting infrastructure issues, such as high CPU usage, low memory, and network latency.

Integration:

- AppDynamics provides extensive integration capabilities, allowing teams to integrate with a wide range of tools and technologies.
- Integrations include popular monitoring tools, such as Splunk and Elasticsearch, as well as IT service management (ITSM) tools like ServiceNow and Remedy.
- AppDynamics also provides REST APIs for custom integrations.

Analytics and Reporting:

- AppDynamics provides analytics and reporting capabilities, allowing teams to track performance metrics and identify areas for improvement.
- Analytics and reporting features include transaction analysis, error analysis, and custom dashboards.
- AppDynamics also provides support for machine learning and predictive analytics, allowing teams to proactively identify and address performance issues.

ArgoCD - Agenda

Session 1: Introduction to ArgoCD

- Overview of ArgoCD and its features
- Understanding the role of ArgoCD in GitOps workflows
- Key concepts and components of ArgoCD

Session 2: Installing and Configuring ArgoCD

- Preparing the environment for ArgoCD installation
- Step-by-step installation guide for ArgoCD
- Configuring ArgoCD server and connecting it to the Git repository

Session 3: ArgoCD Architecture and Components

- Understanding the architecture of ArgoCD
- Exploring the various components of ArgoCD, such as the API server, controller, and repository server

Session 4: Deploying Applications with ArgoCD

- Creating applications in ArgoCD
- Configuring application specifications using GitOps manifests
- Deploying applications and managing their lifecycle with ArgoCD

Session 5: Continuous Delivery with ArgoCD

- Implementing continuous delivery pipelines using ArgoCD
- Automating application updates and rollbacks with ArgoCD
- Monitoring and managing application deployments with ArgoCD

Session 6: Advanced ArgoCD Features

- Exploring advanced features of ArgoCD, such as RBAC and secrets management
- Integrating ArgoCD with other tools and services, like Kubernetes, Helm, and Prometheus

Session 7: Troubleshooting and Best Practices

- Common issues and troubleshooting techniques in ArgoCD
- Best practices for managing and maintaining ArgoCD deployments
- Tips for optimizing performance and scalability in ArgoCD



Threat Model & Tools - STRIDE / PASTA / VAST & Microsoft Threat Modeling Tool / OWASP Threat Dragon

STRIDE:

- STRIDE is a threat modeling methodology developed by Microsoft that identifies six types of threats: Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege.
- The STRIDE approach involves analyzing each component of a system to identify potential threats and vulnerabilities, and then determining appropriate countermeasures to mitigate these risks.

PASTA:

- PASTA is a Process for Attack Simulation and Threat Analysis that is based on the concept of "attacker thinking".
- The PASTA approach involves identifying potential attackers and their motivations, analyzing potential attack paths, and identifying countermeasures to mitigate these risks.
- PASTA is a comprehensive methodology that considers both technical and non-technical factors in the threat modeling process.

VAST:

- VAST is a Visual, Agile, and Simple Threat modeling methodology that is designed to be accessible to both technical and non-technical stakeholders.
- The VAST approach involves creating visual models of the system and its components, and then identifying potential threats and vulnerabilities through a series of brainstorming sessions.
- VAST is designed to be flexible and adaptable, and can be used in a variety of development methodologies.

Microsoft Threat Modeling Tool:

- The Microsoft Threat Modeling Tool is a free tool that helps organizations identify potential threats and vulnerabilities in their software systems.
- The tool uses the STRIDE methodology and allows users to create data flow diagrams of their systems, which can then be analyzed for potential threats and vulnerabilities.
- The Microsoft Threat Modeling Tool provides a comprehensive set of reports and analysis tools, allowing users to prioritize and address potential risks.

OWASP Threat Dragon:

- OWASP Threat Dragon is an open-source tool that helps organizations identify potential threats and vulnerabilities in their software systems.
- The tool uses a data flow diagram approach to threat modeling, allowing users to create visual models of their systems and identify potential attack paths.
- OWASP Threat Dragon provides a comprehensive set of reports and analysis tools, and is designed to integrate with other development tools and methodologies.



Dynamic Application Security Testing (DAST) - OWASP ZAP (Zed Attack Proxy) & Skipfish & Nmap & OpenVAS & Fortify WebInspect

OWASP ZAP (Zed Attack Proxy):

- OWASP ZAP is a free, open-source DAST tool that can be used to identify vulnerabilities in web applications.
- ZAP can be used to perform a variety of tests, including active scanning, passive scanning, and fuzz testing.
- The tool provides a user-friendly interface and can be integrated with other testing tools and frameworks.

Skipfish:

- Skipfish is a free, open-source DAST tool that can be used to identify vulnerabilities in web applications.
- Skipfish is designed to be fast and efficient, making it a good choice for testing large, complex web applications.
- The tool can be run in parallel, allowing it to perform tests on multiple web applications simultaneously.

Nmap:

- Nmap is a free, open-source network scanning tool that can be used to identify open ports and services on a network.
- Nmap can also be used to identify potential vulnerabilities in web applications and other network services.
- The tool provides a variety of scanning options, including stealth scanning and operating system fingerprinting.

OpenVAS:

- OpenVAS is a free, open-source vulnerability scanner that can be used to identify vulnerabilities in web applications and other network services.
- The tool provides a comprehensive set of tests, including active scanning, passive scanning, and vulnerability analysis.
- OpenVAS also provides a variety of reporting options, including detailed reports and risk assessments.

Fortify WebInspect:

- Fortify WebInspect is a commercial DAST tool that can be used to identify vulnerabilities in web applications.
- The tool provides a comprehensive set of tests, including active scanning, passive scanning, and fuzz testing.
- Fortify WebInspect also provides a variety of reporting options, including detailed reports and risk assessments.

Network configurations and Service Discovery - Consul

Network Configurations:

- Consul provides a central service registry that keeps track of all the services in the infrastructure.
- Each Microservices in the infrastructure registers itself with Consul, providing information like its IP address, port, and health status.
- Consul also supports multiple datacenters, allowing for the deployment of services across different regions or availability zones.
- Consul provides a DNS interface that can be used to discover services in the infrastructure. Applications can use this interface to resolve service names to IP addresses and connect to the appropriate service.

Service Discovery:

- Consul provides a service discovery mechanism that enables Microservices to discover and communicate with each other.
- Consul supports different service discovery methods, including DNS, HTTP, and gRPC.
- Consul can perform health checks on the services in the infrastructure to ensure that they are functioning properly. If a service fails a health check, it is removed from the service registry until it is healthy again.
- Consul also supports service segmentation, allowing services to be grouped into logical subsets based on tags or other attributes. This enables more fine-grained control over service discovery and traffic routing.

Software Composition Analysis (SCA) - OWASP Dependency Check & Jfrog Xray

- Introduction to Software Composition Analysis (SCA) and its importance in modern software development.
- Overview of OWASP Dependency Check as a popular SCA tool in the market.
- Understanding the SCA process: scanning, analysis, and remediation.
- Deep-dive into OWASP Dependency Check, including installation, configuration, and usage.
- Demo of OWASP Dependency Check, including a walk-through of its user interface and workflow.
- Understanding OWASP Dependency Check reports and how to interpret them.
- Best practices for using OWASP Dependency Check in SCA, including how to interpret and act on its findings.
- Integration of OWASP Dependency Check with CI/CD pipelines and other development tools.
- Common issues and limitations of SCA tools like OWASP Dependency Check and how to mitigate them.
- Real-world examples of how OWASP Dependency Check has helped organizations improve their software security.
- Future developments in OWASP Dependency Check and SCA in general.
- Q&A session to answer any remaining questions or concerns about OWASP Dependency Check, SCA, or software security in general.
- Introduction to Software Composition Analysis (SCA) and its importance in modern software development.
- Overview of JFrog Xray as a popular SCA tool in the market.
- Understanding the SCA process: scanning, analysis, and remediation.
- Deep-dive into JFrog Xray, including installation, configuration, and usage.
- Demo of JFrog Xray, including a walk-through of its user interface and workflow.
- Understanding JFrog Xray reports and how to interpret them.
- Best practices for using JFrog Xray in SCA, including how to interpret and act on its findings.
- Integration of JFrog Xray with CI/CD pipelines and other development tools.
- Common issues and limitations of SCA tools like JFrog Xray and how to mitigate them.
- Real-world examples of how JFrog Xray has helped organizations improve their software security.
- Future developments in JFrog Xray and SCA in general.
- Comparison between JFrog Xray and other SCA tools in terms of features, capabilities, and pricing.
- Q&A session to answer any remaining questions or concerns about JFrog Xray, SCA, or software security in general.

Runtime application self-protection & Containers (RASP) - Falco & Notary & the Update Framework (TUF) & Nikto

Falco

- Securing Containers (RASP)- Twistkock
- Falco Components
- Userspace program
- Falco Configuration
- Privilege escalation using privileged containers
- Namespace changes using tools like setns
- Read/Writes to well-known directories such as /etc, /usr/bin, /usr/sbin
- Creating symlinks
- Ownership and Mode changes
- Unexpected network connections or socket mutations
- Securing Containers (RASP)- Falco
- Spawned processes using execve
- Falco drivers
- Falco userspace program
- Executing shell binaries such as sh, bash, csh, zsh, etc
- Executing SSH binaries such as ssh, scp, sftp, etc
- Mutating Linux coreutils executables
- Mutating login binaries
- Mutating shadowutil or passwd executables

Notary

- What is CNCF Notary
- Why CNCF Notary?
- What is The Update Framework (TUF)?
- Understand the Notary service architecture
- Brief overview of TUF keys and roles
- Architecture and components
- Example client-server-signer interaction

- Threat model
- Notary server compromise
- Notary signer compromise
- Notary client keys and credentials compromise
- Run a Notary service
- Notary configuration files

Web Application Firewall (WAF) - AWS WAF & Azure Web Application Firewall & Cloudflare Web Application Firewall (WAF)

- Introduction to Web Application Firewall (WAF) and its importance in securing web applications.
- Overview of AWS WAF, Azure Web Application Firewall, and Cloudflare Web Application Firewall as popular WAF solutions in the market.
- Understanding the architecture and key features of each WAF solution.
- Deep-dive into each WAF solution, including installation, configuration, and usage.
- Demo of each WAF solution, including a walk-through of its user interface and workflow.
- Understanding WAF rules and how to create and customize them for specific use cases in each WAF solution.
- Best practices for using each WAF solution to protect web applications, including how to interpret and act on its findings.
- Integration of each WAF solution with other cloud services and development tools.
- Common issues and limitations of WAF solutions like AWS WAF, Azure Web Application Firewall, and Cloudflare Web Application Firewall and how to mitigate them.
- Real-world examples of how each WAF solution has helped organizations improve their web application security.
- Comparison between AWS WAF, Azure Web Application Firewall, and Cloudflare Web Application Firewall in terms of features, capabilities, and pricing.
- Future developments in WAF solutions and WAF in general.
- Q&A session to answer any remaining questions or concerns about AWS WAF, Azure Web Application Firewall, Cloudflare Web Application Firewall, WAF, or web application security in general.

Securing Credentials - HashiCorp Vault & AWS Secrets Manager & Azure Key Vault & AWS KMS & Kubernetes Secrets

- Introduction to securing credentials and why it is important in today's security landscape.
- Overview of HashiCorp Vault, AWS Secrets Manager, Azure Key Vault, AWS KMS, and Kubernetes Secrets as popular solutions for securing credentials.
- Understanding the architecture and key features of each solution.
- Deep-dive into each solution, including installation, configuration, and usage.
- Demo of each solution, including a walk-through of its user interface and workflow.

- Understanding the types of credentials that can be secured using each solution.
- Best practices for using each solution to secure credentials, including how to interpret and act on its findings.
- Integration of each solution with other cloud services and development tools.
- Common issues and limitations of each solution and how to mitigate them.
- Real-world examples of how each solution has helped organizations improve their credential security.
- Comparison between HashiCorp Vault, AWS Secrets Manager, Azure Key Vault, AWS KMS, and Kubernetes Secrets in terms of features, capabilities, and pricing.
- Future developments in securing credentials and the solutions that support them.
- Q&A session to answer any remaining questions or concerns about securing credentials using HashiCorp Vault, AWS Secrets Manager, Azure Key Vault, AWS KMS, Kubernetes Secrets, or credential security in general.

Policy-based control for cloud native environments - Open Policy Agent (OPA)

- Introduction to policy-based control and its importance in cloud native environments.
- Overview of Open Policy Agent (OPA) as a popular open-source policy engine for cloud native environments.
- Understanding the architecture and key features of OPA, including the Rego policy language.
- Deep-dive into OPA, including installation, configuration, and usage.
- Demo of OPA, including a walk-through of its user interface and workflow.
- Understanding how policies can be defined in OPA using the Rego policy language, and how they can be enforced in cloud native environments.
- Best practices for defining policies in OPA to ensure proper control and compliance.
- Integrating OPA with cloud native technologies, such as Kubernetes, Istio, and Envoy, to enforce policies.
- Leveraging OPA to ensure compliance with industry standards and regulations, such as CIS benchmarks and GDPR.
- Monitoring and auditing policy compliance using OPA.
- Common issues and limitations of OPA and how to mitigate them.
- Real-world examples of how OPA has helped organizations improve their policy-based control in cloud native environments.
- Future developments in policy-based control and OPA, including the potential for machine learning and AI-based policy enforcement.
- Comparison between OPA and other policy engines in terms of features, capabilities, and pricing.
- Q&A session to answer any remaining questions or concerns about policy-based control for cloud native environments using Open Policy Agent.

Cloud Security service & Practices - Cloud Security with AWS & Azure service

- Introduction to cloud security and the shared responsibility model for cloud security.
- Overview of AWS security services, such as AWS Identity and Access Management (IAM), AWS Key Management Service (KMS), and AWS Security Hub.
- Best practices for securing AWS resources, including configuring network security, managing access control, and securing data.
- Understanding AWS compliance and regulatory requirements, such as HIPAA and PCI DSS.
- Overview of Azure security services, such as Azure Active Directory (AD), Azure Security Center, and Azure Key Vault.
- Best practices for securing Azure resources, including configuring network security, managing access control, and securing data.
- Understanding Azure compliance and regulatory requirements, such as GDPR and ISO 27001.
- Comparing and contrasting AWS and Azure security services and practices.
- Cloud security automation with tools like AWS CloudFormation and Azure Resource Manager.
- Container security best practices with AWS Elastic Container Service (ECS) and Azure Kubernetes Service (AKS).
- DevSecOps practices and tooling for cloud security.
- Real-world examples of how organizations have successfully implemented cloud security practices with AWS and Azure.
- Future trends and developments in cloud security.
- Q&A session to answer any remaining questions or concerns about cloud security with AWS and Azure.

Security Information and Event Management SIEM - Splunk SIEM

- Introduction to SIEM and its role in security operations.
- Overview of Splunk SIEM, its architecture, and its components.
- Data ingestion and management in Splunk SIEM, including configuring data sources and handling data volume and retention.
- Creating and managing dashboards, reports, and alerts in Splunk SIEM.
- Splunk SIEM search language and syntax, including basic and advanced search commands.
- Using Splunk Enterprise Security (ES) to manage and analyze security events and incidents.
- Integrating third-party security tools and platforms with Splunk SIEM.
- Best practices for deploying and scaling Splunk SIEM in enterprise environments.
- Splunk SIEM use cases and real-world examples of how organizations have successfully implemented it for security operations.
- Future trends and developments in SIEM and Splunk SIEM.
- Q&A session to answer any remaining questions or concerns about Splunk SIEM and SIEM in general.



Thank you!

Connect with us for more info

Call/WhatsApp: - +91 700 483 5930

Mail: - contact@DevOpsSchool.com

www.DevOpsSchool.com























