

Training Course on Productionizing Machine Learning Models with Docker and Kubernetes

Professional Development Training Course Outline

Category	Data Science	Duration	10 Days
Base Fee	\$3,000 USD	Delivery Mode	Online / On-site Hybrid

Course Introduction

Training Course on Productionizing Machine Learning Models with Docker & Kubernetes: Containerization and Orchestration for ML Deployment is meticulously designed to equip data scientists, machine learning engineers, and DevOps professionals with the essential skills and practical knowledge to seamlessly transition machine learning models from development to robust, scalable, and reproducible production environments. By mastering **containerization with Docker** and **orchestration with Kubernetes**, participants will unlock the power of **MLOps best practices**, ensuring efficient **model deployment**, automated **CI/CD pipelines**, and vigilant **model monitoring** for real-world impact. This program is your gateway to becoming a highly sought-after expert in **scalable AI solutions** and **intelligent automation**.

In today's rapidly evolving **AI landscape**, the ability to effectively productionize machine learning models is paramount for organizational success. This course directly addresses the critical challenges of **ML model lifecycle management**, providing hands-on experience with industry-standard tools like **Docker and Kubernetes**. Through practical exercises, real-world **case studies**, and a focus on **end-to-end MLOps workflows**, participants will gain the confidence to build **resilient, high-performance**, and **cost-efficient** ML systems, accelerating innovation and delivering tangible business value.

Programme Curriculum

Training Course on Productionizing Machine Learning Models with Docker & Kubernetes: Containerization and Orchestration for ML Deployment

Introduction

Training Course on Productionizing Machine Learning Models with Docker & Kubernetes: Containerization and Orchestration for ML Deployment is meticulously designed to equip data scientists, machine learning engineers, and DevOps professionals with the essential skills and

practical knowledge to seamlessly transition machine learning models from development to robust, scalable, and reproducible production environments. By mastering **containerization with Docker** and **orchestration with Kubernetes**, participants will unlock the power of **MLOps best practices**, ensuring efficient **model deployment**, automated **CI/CD pipelines**, and vigilant **model monitoring** for real-world impact. This program is your gateway to becoming a highly sought-after expert in **scalable AI solutions** and **intelligent automation**.

In today's rapidly evolving **AI landscape**, the ability to effectively productionize machine learning models is paramount for organizational success. This course directly addresses the critical challenges of **ML model lifecycle management**, providing hands-on experience with industry-standard tools like **Docker and Kubernetes**. Through practical exercises, real-world **case studies**, and a focus on **end-to-end MLOps workflows**, participants will gain the confidence to build **resilient, high-performance**, and **cost-efficient** ML systems, accelerating innovation and delivering tangible business value.

Course Duration

10 days

Course Objectives

1. Deeply understand Docker fundamentals for packaging ML models and their dependencies.
2. Leverage Kubernetes architecture for scalable ML deployments and resource management.
3. Design and implement automated CI/CD pipelines for continuous integration and delivery of ML models.
4. Implement robust strategies for model versioning, experiment tracking, and ensuring reproducible ML workflows.
5. Deploy and manage high-performance ML inference services using Kubernetes.
6. Establish effective model monitoring systems to detect data drift, model decay, and performance degradation.
7. Optimize CPU/GPU allocation and memory sizing for cost-efficient ML workloads on Kubernetes.
8. Apply Kubernetes security fundamentals to protect ML models and sensitive data in production.
9. Develop advanced troubleshooting and debugging skills for production ML systems.
10. Implement strategies for automated model retraining and continuous improvement.
11. Explore concepts of serverless machine learning with Kubernetes.
12. Understand the foundational principles of federated learning and Edge AI in deployment contexts.
13. Discuss best practices for ethical AI deployment and integrating Explainable AI (XAI) principles.

Organizational Benefits

- Rapidly deploy and iterate on ML models, bringing new features and capabilities to users faster.
- Automate repetitive tasks in the ML lifecycle, reducing manual effort and human error.
- Ensure models are continuously monitored, updated, and perform optimally in production.
- Efficiently utilize computing resources, leading to reduced infrastructure costs for ML deployments.
- Foster seamless collaboration between data science, engineering, and operations teams through standardized MLOps practices.
- Guarantee consistency and reproducibility of ML experiments and deployments.

- Build robust and adaptable ML infrastructure capable of handling increasing data volumes and user demands.
- Implement secure deployment practices to protect intellectual property and sensitive data.

Target Audience

1. Data Scientists.
2. Machine Learning Engineers.
3. DevOps Engineers.
4. Software Engineers.
5. Cloud Architects
6. Technical Leads & Managers.
7. AI Practitioners.
8. Solution Architects.

Course Outline

Module 1: Introduction to MLOps and Production ML Challenges

- What is MLOps? Bridging the gap between ML development and operations.
- Common challenges in productionizing ML models (reproducibility, scalability, monitoring).
- Overview of the ML lifecycle in a production context.
- The role of Docker and Kubernetes in solving MLOps challenges.
- **Case Study:** Discussing an organization struggling with inconsistent ML model deployments and slow iteration cycles due to lack of MLOps practices.

Module 2: Docker Fundamentals for ML Models

- Introduction to containers and Docker architecture.
- Creating Dockerfiles for ML environments (Python, dependencies, model artifacts).
- Building and managing Docker images for ML models.
- Docker Compose for multi-container ML applications.
- **Case Study:** Containerizing a fraud detection model and its dependencies for consistent execution across environments.

Module 3: Introduction to Kubernetes for ML Orchestration

- Kubernetes core concepts: Pods, Deployments, Services, Namespaces.
- Setting up a local Kubernetes cluster (Minikube/Kind).
- Deploying a simple containerized ML inference service on Kubernetes.
- Managing Kubernetes resources (YAML definitions).
- **Case Study:** Deploying a simple image classification model as a Kubernetes Deployment and exposing it via a Service.

Module 4: Advanced Docker for ML

- Optimizing Docker images for ML (multi-stage builds, smaller base images).
- Docker registries for storing and sharing ML model images.
- Container security best practices for ML applications.
- Troubleshooting Docker issues in ML contexts.
- **Case Study:** Reducing the size of a large NLP model's Docker image for faster deployment and reduced storage costs.

Module 5: Advanced Kubernetes for ML Deployment

- Scaling ML deployments: Horizontal Pod Autoscaler (HPA) and Vertical Pod Autoscaler (VPA).
- Resource requests and limits for efficient CPU/GPU allocation.
- Kubernetes persistent storage for ML data and models.
- Managing sensitive information with Kubernetes Secrets and ConfigMaps.
- **Case Study:** Auto-scaling a recommendation engine based on real-time user traffic with HPA.

Module 6: CI/CD for Machine Learning with Kubernetes

- Principles of Continuous Integration and Continuous Delivery for ML.
- Integrating Docker and Kubernetes into CI/CD pipelines (e.g., Jenkins, GitLab CI, GitHub Actions).
- Automated testing strategies for ML models in pipelines.
- Implementing GitOps for declarative infrastructure management.
- **Case Study:** Automating the build, test, and deployment process for a new version of a customer churn prediction model.

Module 7: Model Serving Patterns and Strategies

- Understanding different ML model serving patterns (REST APIs, batch inference).
- Implementing RESTful APIs for real-time inference (e.g., FastAPI, Flask).
- Introduction to specialized ML serving frameworks (e.g., KFServing, Seldon Core).
- Load balancing and traffic management for ML services.
- **Case Study:** Deploying a natural language processing (NLP) model with KFServing for efficient inference.

Module 8: Model Monitoring and Observability

- Importance of model monitoring in production (performance, data drift, concept drift).
- Metrics collection for ML models (e.g., Prometheus, Grafana).
- Logging and tracing for debugging production ML systems.
- Alerting mechanisms for anomalies in model behavior.
- **Case Study:** Setting up monitoring for a credit scoring model to detect potential bias or performance degradation over time.

Module 9: MLOps Tools and Ecosystem

- Overview of popular MLOps platforms and tools (e.g., MLflow, Kubeflow).
- Experiment tracking and model registry with MLflow.
- Orchestrating complex ML workflows with Kubeflow Pipelines.
- Comparison of open-source and managed MLOps services.
- **Case Study:** Using MLflow to track experiments and manage model versions for a predictive maintenance project.

Module 10: GPU and Specialized Hardware for ML

- Leveraging GPUs in Kubernetes clusters for deep learning workloads.
- NVIDIA device plugins for GPU allocation.
- Optimizing ML models for GPU inference.
- Considerations for other specialized hardware (TPUs, FPGAs).
- **Case Study:** Deploying a large language model (LLM) on a GPU-enabled Kubernetes cluster for accelerated inference.

Module 11: Data Versioning and Management in MLOps

- Importance of data versioning for reproducible ML.
- Tools and strategies for managing data dependencies (e.g., DVC).
- Data pipelines and feature stores in a production ML context.
- Ensuring data quality and consistency for deployed models.
- **Case Study:** Managing different versions of a dataset used for training a demand forecasting model and ensuring consistency during retraining.

Module 12: Advanced Deployment Strategies & Canary Releases

- Blue/Green deployments for zero-downtime updates.
- Canary deployments for gradual rollout and risk mitigation.
- A/B testing and experimentation with deployed models.
- Rollback strategies for faulty deployments.
- **Case Study:** Performing a canary release for a new version of a recommendation algorithm to a small subset of users before full rollout.

Module 13: Security and Governance in MLOps

- Role-Based Access Control (RBAC) in Kubernetes for ML teams.
- Network policies for isolating ML services.
- Data privacy and compliance considerations for ML deployments.
- Auditing and logging for accountability.
- **Case Study:** Implementing strict RBAC and network policies to secure a financial transaction anomaly detection model.

Module 14: Serverless Machine Learning & Edge AI Concepts

- Introduction to serverless computing for ML inference.
- Deploying ML models with serverless platforms (e.g., Google Cloud Run, AWS Lambda).
- Concepts of Edge AI and deploying models to edge devices.
- Challenges and opportunities of decentralized ML.
- **Case Study:** Discussing the deployment of a small, optimized ML model to IoT devices for real-time anomaly detection at the edge.

Module 15: Future Trends in MLOps and AI

- The rise of Generative AI and its deployment implications.
- Explainable AI (XAI) in production for transparency and trust.
- Federated Learning for privacy-preserving ML.
- MLOps for specialized domains (e.g., MLOps for Reinforcement Learning).
- **Case Study:** Exploring how a company is implementing XAI techniques to explain decisions made by a medical diagnosis AI model.

Training Methodology

This course employs a highly interactive and hands-on training methodology designed for maximum engagement and practical skill acquisition.

- **Instructor-Led Sessions:** Expert-led lectures and discussions on core concepts.
- **Live Coding & Demos:** Real-time demonstrations of Docker, Kubernetes, and MLOps tools.
- **Hands-on Labs:** Extensive practical exercises and coding challenges in a dedicated cloud environment.

- **Case Study Analysis:** In-depth discussion and problem-solving based on real-world ML deployment scenarios.
- **Group Projects:** Collaborative exercises to build and deploy end-to-end ML pipelines.
- **Q&A and Discussion:** Ample opportunities for participants to ask questions and share insights.
- **Best Practices & Pitfalls:** Guidance on industry best practices and common mistakes to avoid.

Register as a group from 3 participants for a Discount

Send us an email: info@fineskilltrainingcenter.com or call +254769199797

Certification

Upon successful completion of this training, participants will be issued with a globally- recognized certificate.

Tailor-Made Course

We also offer tailor-made courses based on your needs.

Key Notes

- The participant must be conversant with English.
- Upon completion of training the participant will be issued with an Authorized Training Certificate
- Course duration is flexible and the contents can be modified to fit any number of days.
- The course fee includes facilitation training materials, 2 coffee breaks, buffet lunch and A Certificate upon successful completion of Training.
- One-year post-training support Consultation and Coaching provided after the course.

Upcoming Scheduled Sessions

Start Date	End Date	Location	Price
07 Sep 2026	18 Sep 2026	Physical	\$3,000
14 Sep 2026	25 Sep 2026	Physical	\$3,000
21 Sep 2026	02 Oct 2026	Physical	\$3,000
28 Sep 2026	09 Oct 2026	Physical	\$3,000

Start Date	End Date	Location	Price
05 Oct 2026	16 Oct 2026	Physical	\$3,000
12 Oct 2026	23 Oct 2026	Physical	\$3,000
19 Oct 2026	30 Oct 2026	Physical	\$3,000
26 Oct 2026	06 Nov 2026	Physical	\$3,000

Ready to enroll or request a customized program? Book online or contact us:

Register Online Now

Email: info@fineskilltrainingcenter.com | Website: www.fineskilltrainingcenter.com