DIGITAL INNOVATION
ACADEMY

AI & ML for DevOps

Explore the practical integration of Machine Learning in DevOps with our comprehensive ML for DevOps course. Designed for DevOps professionals, this course provides the skills needed to apply ML techniques to automate tasks, improve workflows, and enhance operational efficiency.

DAY 1 Foundations of Machine Learning for DevOps

- DevOps meets ML: where, why, and how?
- Observability 2.0: From reactive to predictive
- Case studies: ML for incident prediction, intelligent alerting
- Architecture of ML-enhanced DevOps platforms

Practical ML Integration for DevOps

- Data sources in production (metrics, logs, traces)
- Labeling, model targets, and evaluation strategies for reliability
- ML for SREs: Incident prediction, proactive scaling, root cause hints
- Hands-on lab: build and evaluate an end-to-end ML pipeline for reliability scoring

ML Pipeline for Service Reliability Prediction

- Steps in an ML workflow (data → model → deploy → monitor)
- Feature engineering for system metrics and logs
- Hands-on: Training and evaluating ML models (Random Forest, XGBoost, etc.)
- Interpreting model outputs for operational decision-making

Tools and Environment Setup

- ML environment: Jupyter, Scikit-learn, Pandas, MLflow
- DevOps integrations: Prometheus, Grafana, Loki (overview)
- Version control, experimentation tracking, and containerization
- Dockerizing the ML application for DevOps environments

DAY 2 Time Series Modeling and Forecasting

- Time Series Concepts for DevOps Monitoring
- System metrics as time series (CPU, memory, network, latency)
- Forecasting vs anomaly detection
- Metrics granularity, seasonality, trend, and noise

Anomaly Detection in Logs and Metrics

- Unsupervised approaches: Isolation Forest, DBSCAN
- Statistical vs ML-based detection
- Alerting strategies: thresholding vs intelligent scoring
- Hands-on lab: build an anomaly detection system on real metric/log datasets

Time Series Models and Use Cases

- Classical models: ARIMA, Holt-Winters
- ML/DL-based models: Prophet, LSTM, Facebook Kats
- Comparing model accuracy and stability
- Use case: Forecasting disk usage or latency under load

Practical Deployment for Forecasting Pipelines

- Real-time inference from Prometheus or InfluxDB
- Stream processing with Kafka or Spark Streaming (overview)
- Deploy forecasting API with FastAPI
- Visualize predictions and alerts in Grafana

AI & ML for DevOps

DIGITAL INNOVATION ACADEMY

DAY 3 NLP for Logs and Intelligent Incident Analysis

- Introduction to NLP concepts for DevOps
- Tokenization, embeddings (TF-IDF, Word2Vec, BERT)
- Use cases: log classification, root cause extraction, semantic search
- Building datasets from logs

Building NLP-based Alert Intelligence

- Text classification of log events (error categorization)
- Similarity-based incident grouping
- Use case: GPT-like models for ticket summarization or escalation suggestions
- Hands-on: build and test a basic NLP model for log classification

Capstone Project – Intelligent Monitoring System

- Combine time series, anomaly detection, and NLP modules
- Build a unified intelligent monitoring pipeline
- Evaluate with real or simulated production data
- Peer reviews and solution presentations

Learning outcomes

This training enables DevOps engineers, SREs, and platform teams to leverage machine learning and AI to enhance observability, incident prediction, root cause analysis, and self-healing automation in modern infrastructures.

Your profile

- DevOps engineers
- SRE engineers
- Platform engineers
- Observability engineers
- Technical team leads working on scalable systems and infrastructure.

Prerequisites

- Familiarity with DevOps practices and tooling
- Experience with Python scripting
- Basic understanding of system metrics, logs, and monitoring tools
- Some exposure to Machine Learning is a plus