eMexo TECHNOLOGIES

LEARN, THINK, INNOVATE

MLOps & AIOps

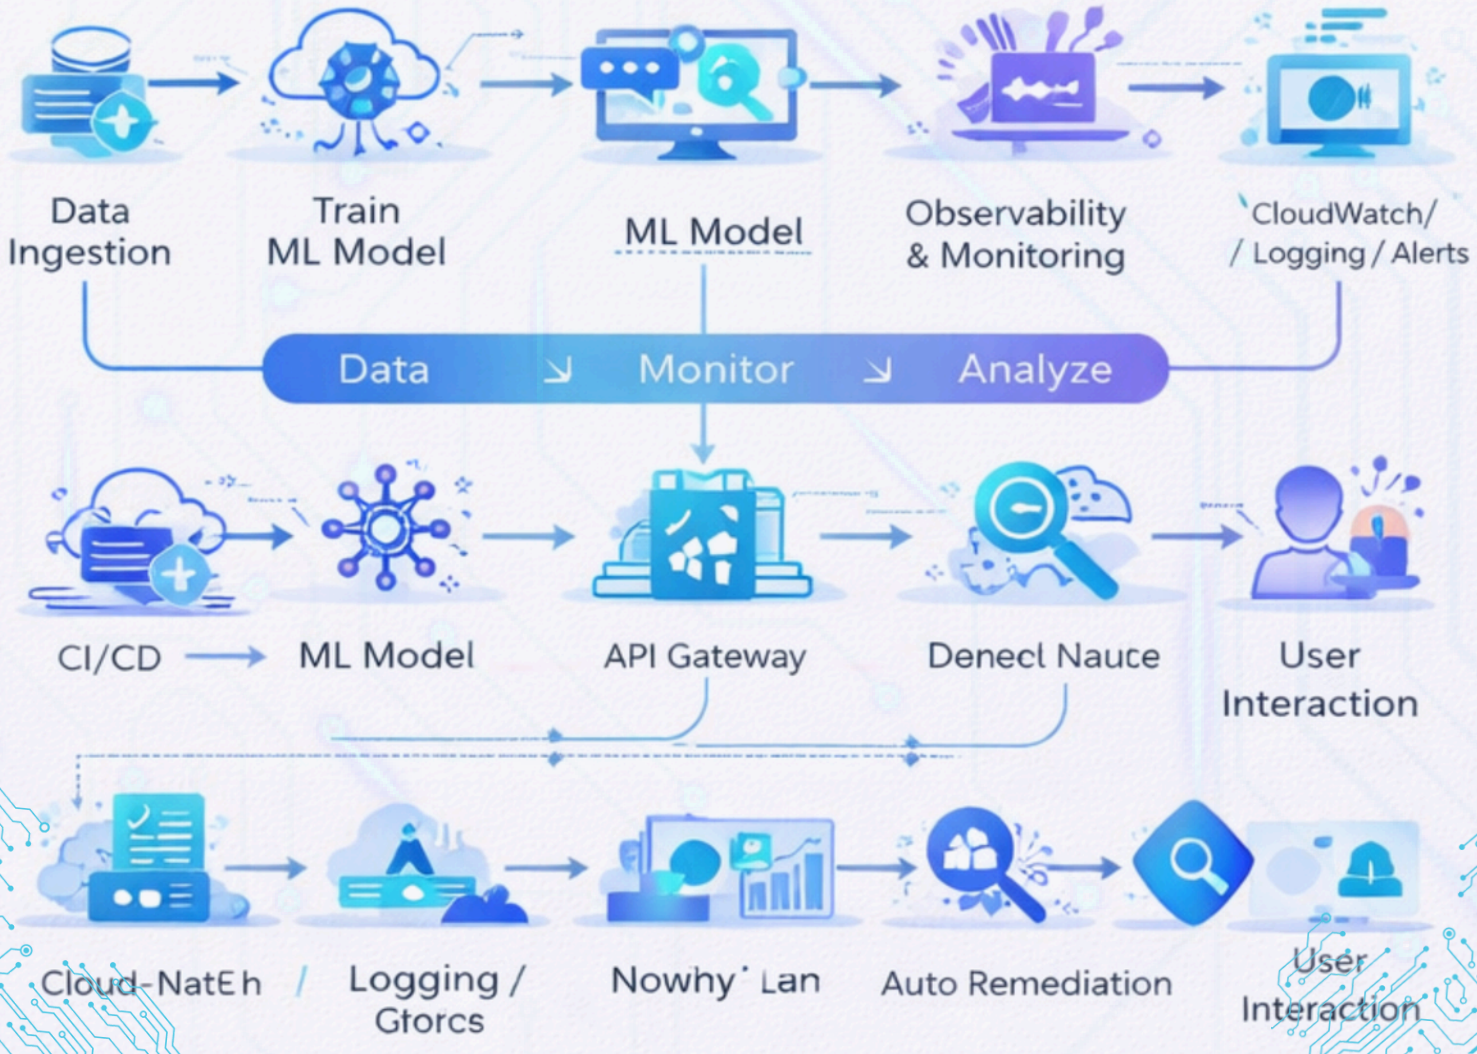

OVERVIEW

The MLOps & AIOps Certification Training is designed for data, machine learning, and AI professionals who want to build, deploy, automate, monitor, and manage intelligent systems at scale in real-world production environments. This program focuses on the complete lifecycle of machine learning models while integrating AIOps practices to enable intelligent monitoring, anomaly detection, root cause analysis, and automated remediation of AI and IT systems.

The course provides a comprehensive understanding of how Machine Learning Operations (MLOps) and Artificial Intelligence for IT Operations (AIOps) work together to ensure reliability, scalability, and efficiency of AI-driven applications. Learners gain practical exposure to automating data pipelines, managing model development and deployment, monitoring model performance and system health, and applying AI-driven insights to operational decision-making.

With strong emphasis on hands-on implementation, the training covers widely used tools and platforms such as MLflow, Kubeflow, Docker, Kubernetes, Git, Jenkins, Airflow, TensorFlow Extended (TFX), along with major cloud services including AWS SageMaker, Azure ML, and Google AI Platform. By the end of the course, participants will be capable of managing production-grade ML systems and AI-powered operations across enterprise environments.

MLOps & AIOps Certification Training – Key Training Features

- Comprehensive coverage of end-to-end ML lifecycle management and AI operations
- Hands-on training with real-world MLOps and AIOps use cases
- Practical implementation of automated data pipelines and ML workflows
- Model versioning, governance, and reproducibility best practices
- CI/CD automation for machine learning and AI systems
- Monitoring of model performance, data drift, and system health

- AIOps-based anomaly detection, root cause analysis, and auto-remediation
- Integration of cloud-native and containerized ML platforms
- Exposure to enterprise-grade tools and frameworks used in production
- Guided learning with structured labs and industry-aligned projects
- Certification-oriented training with practical assessment approach
- Designed for both beginners and experienced professionals

DevOps Fundamentals

Module 1: Introduction to DevOps

- DevOps philosophy and culture
- DevOps lifecycle and practices
- Business value and ROI of DevOps
- DevOps vs traditional development
- DevOps roles and team structures
-

Module 2:Linux Fundamentals for DevOps

- Essential Linux commands and shell scripting
- File system and process management
- User permissions and security basics
- Networking fundamentals
- Automation with bash scripts
-

Module 3:Version Control with Git

- Git fundamentals and workflow
- Branching strategies (GitFlow, trunk-based)
- Pull requests and code reviews
- Git hooks and automation
- GitHub/GitLab features for collaboration
-

Module 4:CI/CD Fundamentals

- Continuous Integration principles
- Continuous Delivery/Deployment concepts
- Pipeline design patterns
- CI/CD tools overview (Jenkins, GitHub Actions)
- Hands-on: Building your first CI/CD pipeline

Module 5:Infrastructure as Code (IaC)

- IaC principles and benefits
- Configuration management with Ansible
- Infrastructure provisioning with Terraform
- Cloud-specific IaC (CloudFormation, ARM templates)
- Hands-on: Automating infrastructure deployment

Module 6:Containerization with Docker

- Container concepts and architecture
- Dockerfile best practices
- Container networking and storage
- Container registries and management
- Hands-on: Containerizing applications

Module 7:Container Orchestration with Kubernetes

- Kubernetes architecture and components
- Deployments, Services, and Pods
- ConfigMaps and Secrets
- Kubernetes networking and storage
- Hands-on: Deploying applications on Kubernetes

Module 8:Cloud Computing Fundamentals

- Major cloud providers comparison (AWS, Azure, GCP)
- Cloud service models (IaaS, PaaS, SaaS)
- Cloud networking and security
- Cost optimization strategies
- Multi-cloud and hybrid cloud approaches

Module 9:Monitoring and Observability

- Monitoring principles and tools
- Prometheus and Grafana setup
- Log aggregation with ELK/EFK stack
- Distributed tracing with Jaeger/Zipkin
- Hands-on: Building a comprehensive monitoring stack

Module 10:DevSecOps Integration

- Security integration in DevOps pipeline
- Vulnerability scanning
- Compliance as code
- Secret management
- Hands-on: Implementing security in CI/CD

MLOps

Module 11:Introduction to MLOps

- ML lifecycle and MLOps concept
- MLOps vs DevOps: Key differences
- MLOps maturity model
- Challenges in ML system deployment
- Industry use cases and success stories

Module 12:Data Engineering for ML

- Data collection and ingestion pipelines
- Data validation and quality checks
- Feature engineering at scale
- Feature stores (Feast, Hopsworks)
- Hands-on: Building data pipelines with Airflow

Module 13:ML Experimentation & Tracking

- Experiment management fundamentals
- Tracking with MLflow and Weights & Biases
- Hyperparameter optimization techniques
- Reproducible ML research
- Hands-on: Setting up experiment tracking

Module 14:ML Version Control

- Versioning ML code with Git
- Data versioning with DVC
- Model versioning strategies
- Experiment reproducibility
- Hands-on: Implementing ML versioning

Module 15:ML Model Packaging & Deployment

- Model serialization formats (ONNX, SavedModel)
- Model serving with RESTful APIs (FastAPI)
- Batch inference systems
- Edge deployment strategies
- Hands-on: Deploying models as services

Module 16:ML CI/CD Pipelines

- ML-specific CI/CD challenges
- Testing ML models and components
- Automating ML workflows
- Continuous training and deployment
- Hands-on: Building an ML CI/CD pipeline

Module 17:Model Monitoring & Management

- Performance monitoring metrics
- Data drift and concept drift detection
- Model retraining strategies
- A/B testing for ML models
- Hands-on: Setting up model monitoring

Module 18:ML Infrastructure Orchestration

- ML workflows with Kubeflow
- Managed ML platforms (SageMaker, Vertex AI)
- Resource optimization for ML
- Scaling ML training and inference
- Hands-on: Orchestrating ML workflows

Module 19:ML on Edge and Mobile

- Edge computing for ML
- Model optimization for edge devices
- TensorFlow Lite and PyTorch Mobile
- Federated learning concepts
- Hands-on: Deploying models to edge devices

Module 20:MLOps for Computer Vision and NLP

- Specific challenges for CV/NLP models
- Data pipeline considerations
- Model optimization techniques
- Deployment architectures
- Hands-on: End-to-end CV/NLP deployment

LLMOps

Module 21:Introduction to LLMOps

- Large Language Model fundamentals
- LLMOps vs traditional MLOps
- LLM lifecycle management
- LLM deployment challenges
- Business applications of LLMs

Module 22:Foundation Model Management

- Open-source vs proprietary LLMs
- Model selection criteria
- Hosting and serving large models
- Model weight management
- Hands-on: Setting up a foundation model

Module 23:Prompt Engineering & Management

- Prompt engineering fundamentals and patterns
- Prompt versioning and templates
- Testing and evaluating prompts
- Prompt management systems
- Hands-on: Building a prompt management workflow

Module 24:LLM Fine-tuning & Customization

- Fine-tuning methodologies
- Parameter-efficient techniques (LoRA, QLoRA)
- Domain adaptation strategies
- Evaluation of fine-tuned models
- Hands-on: Fine-tuning LLMs for specific tasks

Module 25:Retrieval Augmented Generation (RAG)

- RAG architecture and components
- Document retrieval systems
- Vector databases and embeddings
- Hybrid search techniques
- Hands-on: Building a RAG system

Module 26:LLM Deployment Architectures

- Inference optimization techniques
- Quantization and distillation
- Caching strategies
- Scaling and load balancing
- Hands-on: Deploying optimized LLMs

Module 27:LLM Evaluation & Testing

- Evaluation metrics for LLMs
- Red-teaming and adversarial testing
- Automated evaluation frameworks
- Continuous evaluation pipelines
- Hands-on: Building an LLM evaluation system

Module 28:LLM Observability & Monitoring

- Output quality monitoring
- Response time and cost tracking
- User feedback integration
- Anomaly detection for LLMs
- Hands-on: Implementing LLM monitoring

Module 29:Responsible LLM Implementation

- Alignment techniques
- Content filtering systems
- Explainability and transparency
- Ethical considerations and governance
- Hands-on: Implementing LLM guardrails

Module 30:Multimodal LLMs

- Vision-language models
- Audio-text integration
- Multimodal embeddings
- Multimodal fine-tuning strategies
- Hands-on: Working with multimodal LLMs

AI Fundamentals

Module 31:Introduction to AI & Machine Learning

- AI concepts and history
- Types of machine learning
- Deep learning fundamentals
- AI ethics and responsible development
- Current state of AI industry

Module 32:Mathematics for AI

- Linear algebra fundamentals
- Probability and statistics
- Calculus for optimization
- Information theory basics
- Hands-on: Math implementation in Python

Module 33:Machine Learning Fundamentals

- Supervised, unsupervised, and reinforcement learning
- Feature engineering basics
- Model selection and evaluation
- Common ML algorithms
- Hands-on: Building basic ML models

Module 34:Deep Learning Essentials

- Neural network architecture
- Backpropagation and optimization
- Convolutional neural networks
- Recurrent neural networks
- Hands-on: Building deep learning models

Module 35: Transformer Architecture

- Attention mechanisms
- Self-attention and multi-head attention
- Encoder-decoder architecture
- Positional encodings
- Hands-on: Implementing transformer models

Module 36:Natural Language Processing

- Text preprocessing techniques
- Word embeddings and language models
- Sequence modeling for text
- Transformers for NLP
- Hands-on: Building NLP applications

Module 37:Computer Vision Basics

- Image processing fundamentals
- Object detection and recognition
- Image segmentation
- Vision transformers
- Hands-on: Building CV applications

Module 38:Reinforcement Learning

- RL fundamentals and terminology
- Value-based methods
- Policy gradient methods
- Deep reinforcement learning
- Hands-on: Building RL agents

Module 39:AI Tools & Frameworks

- TensorFlow and Keras
- PyTorch ecosystem
- Hugging Face transformers
- JAX for research
- Hands-on: Working with AI frameworks

Module 40:AI Ethics & Governance

- Bias and fairness in AI
- Privacy considerations
- Explainable AI techniques
- Regulatory frameworks
- Hands-on: Implementing ethical AI practices

AIOps

Module 41:Introduction to AIOps

- AIOps concept and evolution
- AIOps vs traditional IT operations
- Business value of AIOps
- AIOps implementation challenges
- AIOps maturity model

Module 42:IT Operations Data Collection

- Telemetry data collection frameworks
- Log aggregation systems
- Metrics collection platforms
- Data integration strategies
- Hands-on: Building data collection pipelines

Module 43:AIOps Data Processing

- Data normalization techniques
- Time-series processing
- Event correlation methods
- Anomaly detection preprocessing
- Hands-on: Processing operations data

Module 44:Anomaly Detection Systems

- Statistical anomaly detection
- ML-based anomaly detection
- Time-series anomaly detection
- Multivariate anomaly detection
- Hands-on: Building anomaly detection models

Module 45:Predictive Analytics for IT

- Failure prediction models
- Capacity planning algorithms
- SLA prediction techniques
- Resource optimization models
- Hands-on: Building predictive models for IT

Module 46:Root Cause Analysis & Remediation

- Automated RCA techniques
- Causal inference in IT systems
- Event correlation for troubleshooting
- Automated remediation frameworks
- Hands-on: Building RCA systems

Module 47:AIOps Platforms & Integration

- Commercial AIOps platforms
- Open-source AIOps tools
- ITSM integration strategies
- Incident management automation
- Hands-on: Implementing an AIOps platform

Module 48:Self-Healing Systems

- Automated remediation patterns
- Self-healing infrastructure
- Chaos engineering practices
- Resilience testing frameworks
- Hands-on: Building self-healing capabilities

Module 49:Cloud-Native AIOps

- Kubernetes observability
- Microservices monitoring
- Serverless function monitoring
- Container health management
- Hands-on: Cloud-native AIOps implementation

Module 50:AIOps & DevSecOps Integration

- Security monitoring with AIOps
- Threat detection models
- Compliance automation
- Security incident response
- Hands-on: Implementing SecOps with AIOps

Generative AI

Module 51:Introduction to Generative AI

- Generative vs discriminative models
- Types of generative models
- Applications of generative AI
- Business use cases
- Ethical considerations

Module 52:Foundation Models

- Pre-training methodologies
- Transfer learning concepts
- Scaling laws and emergent abilities
- Foundation model ecosystems
- Hands-on: Working with foundation models

Module 53:Text Generation Models

- Language model architecture
- GPT and other autoregressive models
- Text generation techniques
- Control mechanisms for text generation
- Hands-on: Building text generation applications

Module 54:Image Generation

- GAN architecture and training
- Diffusion models (DALL-E, Stable Diffusion)
- Text-to-image systems
- Style transfer and image manipulation
- Hands-on: Building image generation applications

Module 55:Audio & Speech Generation

- Speech synthesis technologies
- Music generation models
- Audio style transfer
- Voice cloning considerations
- Hands-on: Building audio generation applications

Module 56:Video Generation

- Text-to-video systems
- Video diffusion models
- Motion synthesis techniques
- Temporal consistency methods
- Hands-on: Building video generation applications

Module 57:Multimodal Generation

- Cross-modal generation techniques
- Text-to-3D systems
- Multimodal understanding
- Combined generative pipelines
- Hands-on: Building multimodal applications

Module 58:Generative AI Deployment

- Serving generative models efficiently
- Latency optimization
- Cost management for generation
- User feedback integration
- Hands-on: Deploying generative AI services

Module 59:Generative AI for Business

- Content creation workflows
- Personalization systems
- Creative assistance tools
- Enterprise integration patterns
- Hands-on: Building business applications

Module 60:Responsible Generative AI

- Bias detection and mitigation
- Content filtering systems
- Attribution and provenance
- Copyright and ownership issues
- Hands-on: Implementing responsible AI guardrails

AI Agents

Module 61: Introduction to AI Agents

- Agent architecture and components
- Types of AI agents
- Agent capabilities and limitations
- Business applications of agents
- Ethical considerations for autonomous systems

Module 62:Agent Development Frameworks

- LangChain for agent development
- AutoGPT architecture
- BabyAGI implementation
- CrewAI for multi-agent systems
- Hands-on: Building your first AI agent

Module 63:Tool Use & Function Calling

- Function calling architecture
- Tool libraries and integration
- API connectivity for agents
- Tool selection reasoning
- Hands-on: Building tool-using agents

Module 64:Agent Memory Systems

- Short-term and working memory
- Long-term knowledge management
- Vector databases for agent memory
- Memory retrieval strategies
- Hands-on: Implementing agent memory

Module 65:Planning & Reasoning

- Planning algorithms for agents
- Chain-of-thought reasoning
- Tree of thought exploration
- Task decomposition techniques
- Goal-oriented behavior
- Hands-on: Building reasoning systems

Module 66:Multi-Agent Systems

- Multi-agent architectures
- Communication protocols
- Role specialization
- Collaborative problem-solving
- Emergent behaviors
- Hands-on: Implementing multi-agent systems

Module 67:Autonomous Decision Making

- Decision theory for agents
- Utility functions and preferences
- Risk assessment and management
- Feedback incorporation
- Hands-on: Building decision-making agents

Module 68:Agent Evaluation & Testing

- Evaluation frameworks for agents
- Benchmarking agent performance
- Simulation environments
- User feedback integration
- Adversarial testing
- Hands-on: Testing agent capabilities

Module 69:Human-Agent Interaction

- Conversational interfaces
- User experience design for agents
- Explainability for agent actions
- Trust building mechanisms
- Hands-on: Designing human-agent interactions

Module 70:Enterprise Agent Deployment

- Agent security considerations
- Scalable agent infrastructure
- Monitoring agent behavior
- Continuous improvement frameworks
- Governance and compliance
- Hands-on: Deploying enterprise-grade agents

** Real-World Projects & Job Preparation **