ClearML ML Engineer Certification

- 3 Days Course
- Hands-on and Lecture

Course Overview

This class prepares students for the ClearML Engineer certification. ClearML is an open-source MLOps platform that enables teams to seamlessly track, orchestrate, and scale machine learning workloads across Kubernetes, cloud, and hybrid environments. By the conclusion of this hands-on training, you will return to work with the skills to deploy, secure, and operate a full ClearML environment — from experiment tracking to GPU-powered model serving.

Throughout the course, you will learn to use Helm, Kubernetes, and cloud-native tools to manage ClearML at scale. You'll configure external data stores, automate agent scaling, integrate with Hugging Face and vLLM, and practice troubleshooting real-world ClearML incidents. The curriculum combines scenario-based labs with production-focused simulations.

Who Should Attend

- Machine Learning Engineers (suggested)
- MLOps Engineers (suggested)
- Data Engineers and DevOps Professionals scaling ML workloads (suggested)

What You'll Learn

- Deploy, secure, and operate a full ClearML environment on Kubernetes and cloud-native infrastructure.
- Automate and orchestrate machine learning pipelines and agent scaling across hybrid environments.
- Configure external data stores and enterprise storage solutions for ClearML.
- Integrate ClearML with open-source tools and LLM serving engines such as Hugging Face and vLLM.
- Troubleshoot production incidents, agent queues, and system performance issues in MLOps workflows.

Outline

1. Core ClearML Architecture & Platform Setup

- Lecture: ClearML Architecture Topology and MLOps Ecosystem Role
- Lecture: Kubernetes Infrastructure Requirements and Ingress Design
- Lecture + Lab: Deploying ClearML using Helm and Custom Values
- Lecture + Lab: Storage Options and External Persistence Configuration

2. Agent Orchestration, Scaling & Pipelines

- Lecture: ClearML Agent Mechanics and Task Queue Governance
- Lecture + Lab: Configuring and Deploying Scalable ClearML Agents
- Lecture + Lab: Dynamic Pipeline Automation and Task Execution

- Lecture + Lab: GPU-Powered Execution and Resource Management

3. Enterprise Integration, Model Serving & Incident Management

- Lecture: Model Serving Integration with Hugging Face and vLLM
- Lecture + Lab: Deploying and Monitoring End-to-End Model Serving Endpoints
- Lecture + Lab: Production Troubleshooting and Agent Queue Debugging
- Lecture + Lab: Capstone Project: End-to-End MLOps Pipeline & Infrastructure Operations

Prerequisites

- Experience with Python programming (suggested)
- Basic familiarity with Linux, Kubernetes, and Helm concepts (suggested)
- Fundamental understanding of Machine Learning and MLOps concepts (suggested)

Next Courses

- ClearML Systems Administrator (CLEARML-SA)
- ClearML Infrastructure Engineer (CLEARML-IE)