



ClearML Systems Administrator

- 3 Days
- Lecture and Hands-on Labs

Course Overview

The ClearML Systems Administrator course at Alta3 Research Labs is a definitive, hands-on training program for Systems Administrators and DevOps Engineers tasked with deploying and securing enterprise-grade MLOps environments. Moving beyond basic usage, students dive deep into the Alta3 lab environment—assisted by VIRGIL, our AI Lab Coach—to master the underlying topology of ClearML, including the complex relationships between Servers, Agents, and Kubernetes. You will learn to architect robust data flows, manage external data stores like MongoDB and Elasticsearch, and handle secrets securely, ensuring you understand not just how to use the platform, but how to own the infrastructure that powers it. A critical focus is placed on transforming standard deployments into secure, multi-tenant ecosystems suitable for large organizations. Participants will master identity integration via OIDC and SSO, implement strict Role-Based Access Control (RBAC), and govern isolated tenants mapped to specific resources. The training culminates in a rigorous Capstone Lab covering “Day-2” operations, where you will deploy Prometheus and Grafana for observability, troubleshoot agent queues, and provision a new tenant end-to-end. By the end of this course, you will possess the specialized skills to architect and maintain a resilient, scalable ClearML platform that satisfies enterprise security and operational standards.

Who Should Attend

- System Administrators

What You'll Learn

- ClearML Architecture Topology
- Secure Multi-Tenancy Implementation
- Identity & Access Governance
- Kubernetes & Agent Orchestration
- Observability & System Monitoring
- Enterprise Operational Troubleshooting

Outline

Core ClearML Platform Operations

- Lecture: The Role of ClearML in the MLOps Ecosystem
- Lecture: Kubernetes Requirements and Ingress Setup
- Lecture + Lab: Install Kubernetes using Ansible
- Lecture + Lab: Persistent Storage for ClearML
- Lecture + Lab: Prepare Ingress Controller for ClearML
- Lecture: ClearML Helm Chart Overview and Customization
- Lecture + Lab: Deploy ClearML via Helm
- Lecture: Connecting External Stores and Managing Secrets

Use ClearML for Tracking & Orchestration

- Lecture: ClearML Architecture: Clients, Server, Agents
- Lecture: Tasks, Projects, and Artifacts Explained
- Lecture + Lab: ClearML Quickstart: Log a Training Run
- Lecture + Lab: Explore Tasks, Projects, and Artifacts in ClearML

Understand ClearML Topology (Stores, Queues, Security)

- Lecture: ClearML Core Components and Data Flow
- Lecture: External Data Stores: MongoDB, Elasticsearch, Redis
- Lecture + Lab: Explore ClearML Architecture Topology
- Lecture: Queues, Agents, and Cache Management
- Lecture + Lab: ClearML Agent: Kubernetes Glue
- Lecture: Security Foundations: TLS, OIDC, RBAC
- Lecture + Lab: Enable TLS and Validate Secure Endpoints

Deploy Tenant Services (Multi-Tenancy & Web Authentication)

- Lecture: Understanding Multi-Tenancy in ClearML
- Lecture: Web Login and Authentication Options (SSO / OIDC / SAML)
- Lecture + Lab: Enable Multi-Tenant Mode via Helm
- Lecture + Lab: Create and Manage Tenants via ClearML SDK
- Lecture + Lab: Map Tenants to Kubernetes Namespaces & Apply Policy

Identity & Authentication (OIDC / SSO / SAML)

- Lecture: Identity Provider Integration: OIDC, SAML, and Authentication Flows
- Lecture: ClearML Authentication Modes: Internal, External, Hybrid
- Lecture + Lab: Deploy and Initialize Keycloak
- Lecture + Lab: Configure ClearML with OIDC (Keycloak / Azure AD Example)
- Lecture: Troubleshooting Authentication & Login Failures

Tenant Administration & Governance

- Lecture: Tenant Governance: RBAC, Groups, Permissions, Separation of Data
- Lecture + Lab: User & Group Lifecycle: Provisioning, Access, Deactivation
- Lecture + Lab: Inspect Tenant Activity, Audit Logs & Usage Patterns
- Lecture: Tenant Lifecycle Best Practices (Provision → Operation → Decommission)

Queue & Agent Operations for Tenants

- Lecture: How ClearML Agents Work
- Lecture: Queue Design: Shared Queues, Tenant Queues, GPU Queues
- Lecture + Lab: Create/Modify Queues for Tenants

Dynamic GPU Fractions - CDMO

- Lecture: Concepts & Mental Models (MIG, MPS, Fractions, Profiles)
- Lecture: Architecture – GPU Operator, Device Plugin, and CDMO Integration
- Lecture: Installation and Version Pinning
- Lecture + Lab: Prepare GPU Nodes for the ClearML Dynamic MIG Operator (CDMO)
- Lecture + Lab: Deploy the NVIDIA GPU Operator
- Lecture: Scheduling, Validation, and Observability
- Lecture + Lab: Enable CDMO for MIG-Based Fractional GPU Scheduling
- Lecture: Troubleshooting and Recovery

ClearML Pipelines

- Lecture + Lab: Initializing Pipelines with the Python SDK
- Lecture + Lab: Pipelines: Build a Simple ML Workflow
- Lecture + Lab: Pipelines: Run Tasks in Parallel
- Lecture + Lab: Pipelines: Reuse Results with Caching
- Lecture + Lab: Pipelines: Execute Remote Workloads from GitHub

Prerequisites

- None