



Vertex AI for Machine Learning Practitioners

Course description

In this comprehensive course, machine learning practitioners will delve into the advanced capabilities of Google Cloud's Vertex AI platform. You will gain a deep understanding of how Vertex AI unifies the entire ML development lifecycle, from data preparation and model training to deployment, monitoring, and MLOps. The course emphasizes practical, hands-on application, enabling you to build, deploy, and manage production-ready machine learning solutions with efficiency and scalability. You will learn to leverage Vertex AI's integrated tools for data management, custom model training, automated hyperparameter tuning, experiment tracking, pipeline orchestration, and responsible AI practices, preparing you to tackle real-world ML challenges effectively.

- Course level: Intermediate
- Duration: 2 days

Activities

This course includes presentations, extensive hands-on labs, detailed demonstrations, and group exercises for collaborative problem-solving.

Course objectives

In this course, you will learn to:

- Summarize the core components and benefits of the Vertex AI platform for ML workflows
- Effectively manage datasets and prepare data for training using Vertex AI tools
- Develop, train, and evaluate custom machine learning models within Vertex AI Workbench and Custom Training
- Optimize model performance through hyperparameter tuning with Vertex AI Vizier and track experiments
- Deploy, monitor, and manage the lifecycle of machine learning models on Vertex AI Endpoints
- Automate and orchestrate end-to-end ML workflows using Vertex AI Pipelines
- Leverage Vertex AI's pre-trained APIs and specialized services for various use cases
- Implement MLOps best practices and responsible AI principles within Vertex AI projects

Intended audience

This course is intended for machine learning practitioners with foundational knowledge of ML concepts and Python:

- Data Scientists
- Machine Learning Engineers
- MLOps Engineers
- AI/ML Researchers
- Solution Architects specializing in ML

Prerequisites

We recommend that attendees of this course have:

- Foundational understanding of machine learning concepts (e.g., supervised/unsupervised learning, model evaluation metrics)
- Proficiency in Python programming
- Basic familiarity with Google Cloud Platform (GCP) or other cloud environments
- Experience with common ML frameworks like scikit-learn, TensorFlow, or PyTorch

Course outline

Module 1: Introducing Vertex AI: The Unified ML Platform

- Overview of Google Cloud's AI/ML landscape and strategic vision
- Deep dive into Vertex AI: its architecture, core services, and benefits for practitioners
- Understanding the MLOps lifecycle and how Vertex AI supports each stage
- Navigating the Vertex AI Console and configuring initial project settings
- Demo: Exploring the Vertex AI console, creating a Vertex AI Workbench instance, and running a simple notebook

Module 2: Data Management and Preparation on Vertex AI

- Managing and versioning datasets using Vertex AI Datasets (tabular, image, video, text)
- Integrating with Google Cloud Storage, BigQuery, and other data sources
- Leveraging Vertex AI Data Labeling for high-quality training data
- Techniques for large-scale data transformation and feature engineering with Dataflow and BigQuery ML
- Hands-on Lab: Ingesting data from BigQuery, creating a managed tabular dataset in Vertex AI, and performing basic data exploration

Module 3: Model Development and Training with Vertex AI

- Developing and iterating on models using Vertex AI Workbench (managed Jupyter notebooks)
- Executing custom training jobs: configuring machine types, accelerators, and distributed training
- Leveraging pre-built training containers and building custom Docker images for complex environments
- Introduction to AutoML Training for tabular, image, and text data for rapid prototyping
- Hands-on Lab: Training a custom scikit-learn model using Vertex AI Custom Training with a pre-built container

Module 4: Experiment Tracking and Hyperparameter Tuning

- Tracking and comparing machine learning experiments with Vertex AI Experiments
- Integrating TensorBoard for advanced visualization of training metrics
- Automating hyperparameter optimization with Vertex AI Vizier (black-box optimization-as-a-service)

- Understanding Vizier's tuning strategies: Grid Search, Random Search, Bayesian Optimization, and algorithms
- Hands-on Lab: Implementing hyperparameter tuning for a TensorFlow model using Vertex AI Vizier and tracking results in Vertex AI Experiments

Module 5: Model Deployment, Prediction, and Monitoring

- Deploying trained models to Vertex AI Endpoints for online predictions with autoscaling
- Performing batch predictions for large, offline inference tasks
- Configuring model monitoring for drift detection (data drift, concept drift) and data quality issues
- Understanding and applying Vertex AI Explainable AI (feature attributions, counterfactual explanations)
- Hands-on Lab: Deploying a custom-trained model to a Vertex AI Endpoint, making online predictions, and setting up model monitoring

Module 6: Automating ML Workflows with Vertex AI Pipelines

- Introduction to MLOps principles and the critical role of automated pipelines
- Designing and orchestrating end-to-end ML workflows using Vertex AI Pipelines (Kubeflow Pipelines SDK)
- Developing custom pipeline components and reusing existing ones
- Integrating various Vertex AI services and other GCP services within pipeline steps
- Hands-on Lab: Building and executing a multi-step ML pipeline on Vertex AI, including data preparation, training, and model registration

Module 7: Leveraging Pre-trained Models and Specialized Services

- Overview of Vertex AI pre-trained APIs: Vision AI, Natural Language AI, Translation AI, Speech-to-Text, Text-to-Speech
- Customizing pre-trained models with Vertex AI's advanced transfer learning capabilities
- Introduction to Vertex AI Search and Vertex AI Conversation for enterprise search and conversational AI solutions
- Exploring Vector Search for semantic similarity and retrieval-augmented generation (RAG) applications
- Demo: Interacting with Vertex AI Vision API for image analysis and custom model deployment

Module 8: Advanced MLOps Practices and Responsible AI

- Implementing robust MLOps best practices: CI/CD for ML, model versioning, lineage tracking, and governance
- Introduction to Responsible AI on Vertex AI: fairness, interpretability, privacy, and safety considerations
- Securing Vertex AI resources with Identity and Access Management (IAM) and network configurations
- Strategies for cost management and optimization within Vertex AI projects
- Course Wrap-up: Next steps for advanced Vertex AI usage, exploring Google Cloud certifications, and additional learning resources

Your course overview ends here.
But your cloud journey doesn't!

Unlock advanced training and certification pathways with **CloudThat**.



About CloudThat

Established in 2012, CloudThat is an award-winning company and the first in India to offer cloud training and consulting services for individuals and enterprises worldwide. Recently, it won Google Cloud's New Training Partner of the Year Award for 2025, becoming the first company in the world in 2025 to hold awards from all three major cloud giants: AWS, Microsoft, and Google. CloudThat notably won consecutive AWS Training Partner of the Year (APJ) awards in 2023 and 2024 and the Microsoft Training Services Partner of the Year Award in 2024, bringing its total award count to an impressive 12 awards in the last 8 years.

CloudThat specializes in Cloud Migration, Data Platforms, DevOps, IoT, and cutting-edge technologies like Gen AI & AI/ML. As a Microsoft Solutions Partner, AWS Premier Tier Training Partner, Google Cloud Platform Partner, and collaborator with leading organizations like HPE, Oracle, Nvidia and Databricks, CloudThat has trained over 850,000 professionals across 600+ cloud certifications, empowering students and professionals worldwide to advance their skills and careers.

To know more about our cloud certification training, email at sales@cloudthat.com or call us at [+91-8880002200](tel:+91-8880002200) or [+1-8555588830](tel:+1-8555588830) for those from outside India.