# AI Platform Pipelines documentation

**Beta**

This product or feature is covered by the Pre-GA Offerings Terms (/terms/service-terms#1) of the Google Cloud Platform Terms of Service. Pre-GA products and features may have limited support, and changes to pre-GA products and features may not be compatible with other pre-GA versions. For more information, see the launch stage descriptions (/products#product-launch-stages).

Machine learning (ML) workflows include steps to prepare and analyze data, train and evaluate models, deploy trained models to production, track ML artifacts and understand their dependencies, etc. Managing these steps in an ad-hoc manner can be difficult and time-consuming.

*MLOps* is the practice of applying DevOps practices to help automate, manage, and audit ML workflows. AI Platform Pipelines helps you implement MLOps by providing a platform where you can orchestrate the steps in your workflow as a pipeline. ML pipelines are portable and reproducible definitions of ML workflows.

AI Platform Pipelines makes it easier to get started with MLOps by saving you the difficulty of setting up Kubeflow Pipelines with TensorFlow Extended (TFX). Kubeflow Pipelines (https://www.kubeflow.org/docs/pipelines/overview/pipelines-overview/) is an open source platform for running, monitoring, auditing, and managing ML pipelines on Kubernetes. TFX (https://www.tensorflow.org/tfx) is an open source project for building ML pipelines that orchestrate end-to-end ML workflows.

## Overview

### Introduction to AI Platform Pipelines

An overview of AI Platform Pipelines.
(/ai-platform/pipelines/docs/introduction)

## Getting started

### Getting started with AI Platform Pipelines

Quickstart guide to setting up AI Platform Pipelines and running a pipeline.

(/ai-platform/pipelines/docs/getting-started)

## How-to guides

### Setting up AI Platform Pipelines

Learn how to set up AI Platform Pipelines.

(/ai-platform/pipelines/docs/setting-up)

### Creating an ML pipeline

Learn how to orchestrate your ML process as a pipeline.

(/ai-platform/pipelines/docs/create-pipeline)

### Running an ML pipeline

Learn how to access the Kubeflow Pipelines dashboard and run pipelines.

(/ai-platform/pipelines/docs/run-pipeline)

### Connecting to AI Platform Pipelines using the Kubeflow Pipelines SDK

Learn how to connect to your AI Platform Pipelines cluster using the Kubeflow Pipelines SDK.

(/ai-platform/pipelines/docs/connecting-with-sdk)

### Configuring your GKE cluster

Configure your Google Kubernetes Engine cluster to ensure that AI Platform Pipelines has sufficient computational resources and access to Google Cloud resources, such as Cloud Storage or BigQuery.
(/ai-platform/pipelines/docs/configure-gke-cluster)

## Support

### Troubleshooting
Troubleshooting common problems.
(/ai-platform/pipelines/docs/troubleshooting)

### Getting support
Get assistance with AI Platform Pipelines issues.
(/ai-platform/pipelines/docs/getting-support)

### Billing questions
Get answers to common billing questions.
(/ai-platform/pipelines/docs/billing-questions)

## Resources

### Pricing
Learn about AI Platform Pipelines pricing.
(/ai-platform/pipelines/pricing)

### Release notes

Learn about updates to AI Platform Pipelines.
(/ai-platform/pipelines/docs/release-notes)