


Course Summary

Module 6

Google Cloud Big Data and Machine Learning Fundamentals



Congratulations!




Google Cloud Big Data and Machine Learning Fundamentals

Google Cloud

Congratulations on completing the Big Data and Machine Learning Fundamentals course! We hope that you learned some valuable information from this course that will help advance your career.

Data-to-AI lifecycle



Throughout the course, we introduced a number of products and technologies to support Google's **data-to-AI lifecycle**.

01 Big Data and Machine Learning on Google Cloud ☒

02 Data Engineering for Streaming Data ☐

03 Big Data with BigQuery ☐

04 Machine Learning Options on Google Cloud ☐

05 The Machine Learning Workflow with Vertex AI ☐

Let's do a final review of the main concepts presented. **In the first module** of the course, you were introduced to the Google Cloud infrastructure and Google's big data and machine learning products.

The Google Cloud architecture



Of the three layers of the Google Cloud infrastructure, you explored the middle and top layers.

- On the middle layer sit **compute** and **storage**. Google Cloud decouples compute and storage so they can scale independently based on need.
- And on the top layer sit the **big data and machine learning products**, which enable you to perform tasks to ingest, store, process, and deliver business insights, data pipelines, and machine learning models.

01 Big Data and Machine Learning on Google Cloud ☒

02 Data Engineering for Streaming Data ☒


03 Big Data with BigQuery ☐

04 Machine Learning Options on Google Cloud ☐

05 The Machine Learning Workflow with Vertex AI ☐

In the **second module of the course**, you explored data engineering for streaming data.

Streaming data pipeline



This included how to build a streaming data pipeline—from **ingestion** with **Pub/Sub**, to **processing** with **Dataflow**, and finally, to **visualization** using **Looker** and **Looker Studio**.

01 Big Data and Machine Learning on Google Cloud ☒

02 Data Engineering for Streaming Data ☒

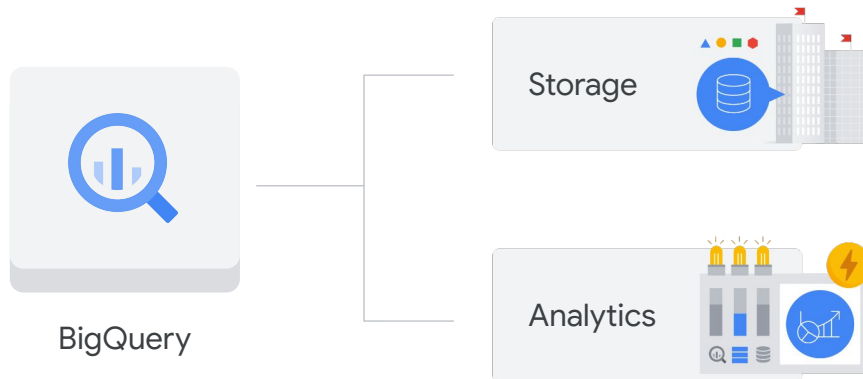
03 Big Data with BigQuery ☒

04 Machine Learning Options on Google Cloud ☐

05 The Machine Learning Workflow with Vertex AI ☐


After that, in the **third module of the course**, you were introduced to BigQuery, which is Google's fully-managed data warehouse.

BigQuery



BigQuery provides two services in one: storage *plus* analytics.

BigQuery ML



You also learned about BigQuery ML, the machine learning tool used for developing machine learning models directly in BigQuery.

01 Big Data and Machine Learning on Google Cloud ☒

02 Data Engineering for Streaming Data ☒

03 Big Data with BigQuery ☒


04 Machine Learning Options on Google Cloud ☒

05 The Machine Learning Workflow with Vertex AI ☐

In the **fourth module of the course**, you explored the options available to build and deploy machine learning models with Google Cloud.

ML options on Google Cloud

Familiar with SQL and have data in BigQuery



01

BigQuery ML

Have little ML expertise



02

Pre-built APIs

Want to build custom models with your own training data with minimal coding



03

AutoML

Want full control of the ML workflow



04

Custom training

Pre-built
containers

Custom
containers

Google Cloud

- If you're familiar with SQL and already have data in BigQuery, you can use **BigQuery ML** to develop machine learning models.
- If you have little ML experience, using **pre-built APIs** is likely the best choice. Pre-built APIs address common perceptual tasks such as vision, video, and natural language. They're ready to use without any ML expertise or model development effort.
- If you want to build custom models with your own training data while spending minimal time coding, then **AutoML** is a great choice. AutoML provides a code-less solution to enable you to focus on business problems instead of the underlying model architecture and ML provisioning.
- And if you want full control of the machine learning workflow, Vertex AI **custom training** lets you train and serve custom models with code on **Vertex Workbench**.
 - Using **pre-built containers**, you can leverage popular ML libraries, such as Tensorflow and PyTorch.
 - Alternatively, you can build a **custom container** from scratch.

01 Big Data and Machine Learning on Google Cloud ☒

02 Data Engineering for Streaming Data ☒


03 Big Data with BigQuery ☒

04 Machine Learning Options on Google Cloud ☒

05 The Machine Learning Workflow with Vertex AI ☒

In the **final module of the course**, you learned about the machine learning workflow using Vertex AI, a unified platform that brings all the components of the machine learning ecosystem and workflow together.


The three stages of the ML workflow



The machine learning workflow comprises three stages.

- In stage one, data preparation, data is uploaded and feature engineering is applied.
- In stage two, model training, the model is trained and evaluated.
- And in stage three, model serving, the model is deployed and monitored.


cloud.google.com/training/data-engineering-and-analytics
cloud.google.com/training/machinelearning-ai



We hope that this course is just the beginning of your big data and machine learning journey.

For more training and hands-on practice with data engineering and analytics, please refer to cloud.google.com/training/data-engineering-and-analytics. And for more training with machine learning and AI, please explore the options available at cloud.google.com/training/machinelearning-ai.

cloud.google.com/certifications



Google Cloud certification

Google Cloud

And if you're interested in validating your expertise and showcasing your ability to transform businesses with Google Cloud technology, you might consider working toward a Google Cloud certification.

You can learn more about Google Cloud's certification offerings at cloud.google.com/certifications.



Thanks for completing this course. We'll see you next time!