

Training: Google Cloud

Data Warehousing with BigQuery: Storage Design, Query Optimization, and Administration



TRAINING GOALS:

In this course, you learn about the internals of BigQuery and best practices for designing, optimizing, and administering your data warehouse. Through a combination of lectures, demos, and labs, you learn about BigQuery architecture and how to design optimal storage and schemas for data ingestion and changes. Next, you learn techniques to improve read performance, optimize queries, manage workloads, and use logging and monitoring tools. You also learn about the different pricing models. Finally, you learn various methods to secure data, automate workloads, and build machine learning models with BigQuery ML.

What you'll learn

- Describe BigQuery architecture fundamentals.
- Implement storage and schema design patterns to improve performance.
- Use DML and schedule data transfers to ingest data.
- Apply best practices to improve read efficiency and optimize query performance.
- Manage capacity and automate workloads.
- Understand patterns versus anti-patterns to optimize queries and improve read performance.
- Use logging and monitoring tools to understand and optimize usage patterns.
- Apply security best practices to govern data and resources.
- Build and deploy several categories of machine learning models with BigQuery ML

Audience

Data analysts, data scientists, data engineers, and developers who perform work on a scale that requires advanced BigQuery internals knowledge to optimize performance.

CONSPECT:

- BigQuery Architecture Fundamentals
 - Introduction
 - BigQuery Core Infrastructure
 - BigQuery Storage
 - BigQuery Query Processing

- BigQuery Data Shuffling
- Storage and Schema Optimizations
 - BigQuery Storage
 - Partitioning and Clustering
 - Nested and Repeated Fields
 - ARRAY and STRUCT syntax
 - Best Practices
- Ingesting Data
 - Data Ingestion Options
 - Batch Ingestion
 - Streaming Ingestion
 - Legacy Streaming API
 - BigQuery Storage Write API
 - Query Materialization
 - Query External Data Sources
 - Data Transfer Service
- Changing Data
 - Managing Change in Data Warehouses
 - Handling Slowly Changing Dimensions (SCD)
 - DML statements
 - DML Best Practices and Common Issues
- Improving Read Performance
 - BigQuery's Cache
 - Materialized Views
 - BI Engine
 - High Throughput Reads
 - BigQuery Storage Read API
- Optimizing and Troubleshooting Queries
 - Simple Query Execution
 - SELECTs and Aggregation
 - JOINs and Skewed JOINs
 - Filtering and Ordering
 - Best Practices for Functions
- Workload Management and Pricing
 - BigQuery Slots
 - Pricing Models and Estimates

- Slot Reservations
- Controlling Costs
- Logging and Monitoring
 - Cloud Monitoring
 - BigQuery Admin Panel
 - Cloud Audit Logs
 - INFORMATION_SCHEMA
 - Query Path and Common Errors
- Security in BigQuery
 - Secure Resources with IAM
 - Authorized Views
 - Secure Data with Classification
 - Encryption
 - Data Discovery and Governance
- Automating Workloads
 - Scheduling Queries
 - Scripting
 - Stored Procedures
 - Integration with Big Data Products
- Machine Learning in BigQuery
 - Introduction to BigQuery ML
 - How to Make Predictions with BigQuery ML
 - How to Build and Deploy a Recommendation System with BigQuery ML
 - How to Build and Deploy a Demand Forecasting Solution with BigQuery ML
 - Time-Series Models with BigQuery ML
 - BigQuery ML Explainability

REQUIREMENTS:

Big Data and Machine Learning Fundamentals

Difficulty level



CERTIFICATE:

The participants will obtain certificates signed by Google Cloud.

TRAINER:

Authorized Google Cloud Trainer.