

Understanding Cloud Spanner

In this course, you learn about Cloud Spanner. You will get an introduction to Cloud Spanner and understand how it differs from other database products. You also learn when and how to use Cloud Spanner to solve your relational database needs at scale.

**DURATION**

3 days

**LEVEL**

200 - Intermediate

**FORMAT**

Instructor-led

What you'll learn

- Build scalable, managed, relational databases by using Cloud Spanner.
- Create and manage Cloud Spanner databases by using the CLI, Terraform, Python API, and the Google Cloud console.
- Program and run queries and transactions by using the Cloud Spanner API
- Integrate Cloud Spanner with applications.

Overview	6 modules · 10 labs
Who this course is for	• Customers
Products	• Spanner • Spanner Studio • Confidential Cloud Run • Dataflow Prime • Cloud Spanner • Cloud Functions • Cloud Run • Dataflow
Prerequisites	• Some prior Google Cloud experience at the fundamental level and experience with relational databases, the SQL language, and some programming are assumed.

Module 01

The Need for Spanner

Topics	• What is Spanner? • Spanner and the CAP Theorem • History of Spanner • Cloud Spanner Use Cases
Objectives	• Explain the core concepts and features of Cloud Spanner. • Understand how Cloud Spanner fits in the CAP theorem. • Describe the history of Cloud Spanner. • Explain Cloud Spanner use cases.

Module 02

Getting Started with Spanner

Topics

- Planning Spanner Instances
- Automating Instance Creating
- Creating Databases in Spanner

Objectives

- Architect Cloud Spanner instances based on location, capacity, availability, and cost.
- Create Spanner instances by using the Google Cloud console, Google Cloud CLI, and Terraform.
- Create Spanner databases by using SQL

Module 03

Optimizing Spanner Schemas

Topics

- Spanner Architecture
- Choosing Primary Keys
- Defining Database Schemas in Spanner
- Understanding Interleaving and Foreign Keys
- Understanding Secondary Indexes

Objectives

- Optimize schemas for Spanner architecture.
- Choose appropriate primary keys.
- Manage relationships with primary and foreign keys and with interleaved tables.

Module 04

Programming Spanner Applications, Queries, and Transactions

Topics

- Authentication and Authorization
- Using the Spanner Client Libraries
- Running Queries
- Managing Transactions

Objectives

- Authenticate users and applications that access Spanner databases using Identity Access Management.
- Program Spanner applications using Google Cloud client libraries and Python.
- Optimize queries using strong reads, stale reads, and indexes.
- Manage transactions in Spanner.

Module 05

Deploying Spanner Applications

Topics	• Using Spanner from Applications • Building Data Pipelines into and out of Spanner
Objectives	• Deploy Spanner applications to Google Cloud serverless runtimes. • Migrate data to and from Cloud Spanner by using Dataflow jobs and Apache Beam.
Activities	• Lab: Deploying Spanner Applications with Cloud Functions and Cloud Run

Module 06

Spanner Administration

Topics	• Managing your Data in Spanner • Managing Change Operations
Objectives	• Administer Cloud Spanner instances. • Backup, restore, import, and export data. • Modify database schemas with no downtime. • Monitor your Cloud Spanner databases and applications
Activities	• Lab: Reconciling Account Data with Cloud Spanner Change Streams • Lab: Leverage the Autoscaler Tool for Cloud Spanner to Achieve Workload Elasticity

Module 07

Spanner Best Practices

Topics	• Spanner Best Practices • Challenge Lab
Objectives	• Review best practices for using Cloud Spanner
Activities	• Challenge Lab: Administering a Spanner Database