

DP-203: Data Engineering in Microsoft Azure, Part 2 of 7: Relational Data Stores

page 1

Meet the expert: Eshant Garg has 16 years of extensive professional experience with expertise in Database and Business Intelligence Solutions, Advanced Analytics, Design and Solution Architect, Reporting, and Cloud Computing Technologies (Azure & AWS).

As a developer and architect, he has worked closely with customers, users, and colleagues to support business solutions across a variety of industries including healthcare, insurance, finance, and government ranging from small companies to fortune 500 companies.

Prerequisites: AZ-900 Azure Fundamentals is very helpful but not required. A Candidate for the exam must have strong knowledge of data processing languages such as SQL, Python, or Scala, and they need to understand parallel processing and data architecture patterns.

Runtime: 06:18:48

Course description: The DP-203 Exam is measured in Four domains: Design and implement data storage (40-45%), Design and develop data processing (25-30%), Design and implement data security (10-15%), and Monitor and optimize data storage and data processing (10-15%).

This course covers Data Storage using Relational Data Stores.

Course outline:

Implement Relational Data

- Introduction
- Learning Objectives
- Why Choosing Sql Server In Azure
- Azure Laas Vs Pass Database Offerings
- Sql Server Paas Deployment Options
- Demo: Single Database
- Summary

DB Purchasing Models and Service Tiers

- Introduction
- Purchasing Models And Service Tiers
- Elastic Db
- Demo: Elastic DB
- Summary

Security

- Introduction
- Security
- Scaling Azure DB
- Sql Db HA and DA Options
- Summary

SQL DB Backup and Restore

- Introduction
- Sql Db Backup And Restore
- Traditional Vs Modern Warehouse Architecture
- Summary

What is Synapse Analytics Service

- Introduction
- What Is Synapse Analytics Service

- Create Dedicated Sql Pool
- Connect Sql Pool With Ssms
- Create Azure Synapse Analytics Workspace
- Summary

Explore Synapse Studio

- Introduction
- Synapse Studio
- Create Dedicated Sql Pool And Spark Pool From Ins
- Analyse Data Using Dedicated Sql Pool
- Summary

Analyze Data using Apache Spark Notebook

- Introduction
- Analyse Data Using Apache Spark Notebook
- Analyse Data Using Serverless SQL Pool
- Data Factory Copy Tool From Synapse Integrate Tab
- Summary

Azure Synapse Architecture

- Introduction
- Azure Synapse Architecture
- Storage And Sharding Patterns
- Data Distribution And Distributing Keys
- Data Types And Table Types
- Partitioning
- Summary

Best Practices for Fact and Dimension tables

- Introduction
- Best Practices For Fact And Dimension Tables
- Demo: Analyse Data Distribution Before Migration

- Different Loading Methods
- Loading With Ssis Vs Polybase
- Summary

Loading with Polybase

- Introduction
- Demo: Loading With Polybase
- Polybase Setup
- Summary

Scaling Azure DW

- Introduction
- Scaling Azure Dw
- Sql Dw Backup And Restore
- Azure DB vs. Dw
- Summary

Dynamic Data Masking

- Introduction
- Dynamic Data Masking
- Dynamic Data Description
- Summary

Encryption

- Introduction
- Encryption
- Encryption In Rest
- Summary