



Implementing Microsoft Azure Cosmos DB Solutions

Course 20777: 2 days; Instructor-Led

Introduction

This two-day instructor-led course is aimed at database professionals who are looking to implement a Cosmos DB solution.

At Course Completion

After completing this course, students will be able to:

- Describe the purpose and architecture of Azure Cosmos DB.
- Describe how to design documents and collections to meet business requirements, and how to use the SQL API to build applications that use these documents.
- Describe how to create user-defined functions, stored procedures, and triggers.
- Describe how to tune a database, and how to monitor performance.
- Describe how to create efficient Graph database models using Cosmos DB.
- Describe how to use Azure Search, HDInsight, Azure Databricks, and Power BI with Cosmos DB to query and analyze big data.
- Describe how to use Cosmos DB as a source and sink for streaming data.

Audience

The primary audience for this course is database developers and architects (IT professionals, developers, and information workers) who plan to implement big data solutions on Azure using Cosmos DB.

Prerequisites

In addition to their professional experience, students who attend this training should already have the following technical knowledge:

- The fundamental concepts of partitioning, replication, and resource governance for building and configuring scalable applications that are agnostic of a Cosmos DB API.
- A basic working knowledge of the Cosmos DB SQL API.

Course Outline

Module 1: Introduction to Azure Cosmos DB

Describe the purpose and architecture of Cosmos DB.

Lessons

- Review of NoSQL database structures
- Migrating data and applications to Cosmos DB
- Managing data in Cosmos DB

Lab : Creating and using a SQL API database in Cosmos DB

- Creating and configuring a Cosmos DB database
- Migrating data from a Mongo DB database to Cosmos DB
- Using the SQL API to access data
- Protecting data in a Cosmos DB database

After completing this module, students will be able to:

- Create and configure a Cosmos DB.
- Migrate data from a Mongo DB database to a Cosmos DB database.
- Describe accessing data using the SQL API.
- Describe how to protect data in a Cosmos DB.

Module 2: Designing and Implementing SQL API Database Applications

This module describes how to design documents and collections to meet business requirements, and how to use the SQL API to build applications that use these documents.

Lessons

- Document models in Cosmos DB
- Querying data in a SQL API database
- Querying and maintaining data programmatically

Lab : Designing and implementing SQL API database applications

- Design the document structure & partitioning strategy for the product catalog for the retail system
- Importing product catalog data
- Querying product catalog information
- Maintaining stock levels in the product catalog

After completing this module, students will be able to:

- Design NoSQL document structures that support business requirements and enable efficient operations.
- Describe how to perform SQL queries against a SQL API database.
- Explain how to insert, modify, delete, and query data in a SQL API database programmatically.

Module 3: Implementing Server Side Operations

Describe how to create user-defined functions, stored procedures, and triggers.

Lessons

- Server-side programming with Cosmos DB
- Creating and using stored procedures
- Using triggers to maintain data integrity

Lab : Writing user-defined functions, stored procedures and triggers

- Design and implement the document and collection structure
- Implement the shopping cart functionality in the online retail system.
- Extend the online retail system to create orders from the items in a shopping cart.
- Extend the online retail system further to enable customers to view orders and backorders.

After completing this module, students will be able to:

- Describe how Cosmos DB enables you to implement server-side operations by writing JavaScript code.
- Describe how use the JavaScript Query API to implement transactional operations with stored procedures.
- Describe how use create triggers that you can use to maintain integrity when inserting, updating, and deleting documents.

Module 4: Optimizing and monitoring performance

Describe how to optimize a database, and how to monitor performance.

Lessons

- Optimizing database performance
- Monitoring the performance of a database

Lab : Tuning a database and monitoring performance

- Gathering execution statistics
- Examining how the different consistency models can impact throughput and latency
- Investigate the effects of triggers on performance
- Monitoring performance and tuning the partition key

After completing this module, students will be able to:

- Describe how to tune the configuration of a database and collections for best performance.
- Describe how to assess the performance of a document database and identify options for improving throughput.

Module 5: Designing and Implementing a Graph Database

This module describes how to create efficient graph database models using Cosmos DB.

Lessons

- Graph database models in Cosmos DB
- Designing Graph database models for efficient operation

Lab : Designing and implementing a Graph database

- Implementing a recommendations engine for customers
- Recording product purchase information
- Query a Graph database to obtain analytics

After completing this module, students will be able to:

- Describe the features that Cosmos DB provides for implementing graph databases.
- Design NoSQL graph structures that support business requirements and enable efficient operations.

Module 6: Querying and Analyzing Big Data with Cosmos DB

This module describes how to use Azure Search and HDInsight with Cosmos DB to query and analyze big data.

Lessons

- Integrating Cosmos DB with Azure search to optimize queries
- Analyzing data in a Cosmos DB database using Apache Spark
- Visualizing data in a Cosmos DB database

Lab : Querying and Analyzing Big Data with Cosmos DB

- Extending product search capabilities
- Performing end-of-month processing
- Visualizing sales data
- Exploring sales data

At the end of this module, students will be able to:

- Describe how to integrate Cosmos DB with Azure Search to perform efficient query processing over big data.
- Describe how to analyze big data held in Cosmos DB using Apache Spark.
- Describe how to visualize data in Cosmos DB using Jupyter notebooks, Power BI, and Azure Databricks.

Module 7: Implementing Stream Processing with Cosmos DB

This module describes how to use Cosmos DB as a source and sink for streaming data.

Lessons

- Working with the Cosmos DB change feed
- Integrating Cosmos DB into streaming solutions
- Lab : Using Cosmos DB with stream processing
- Handling orders
- Maintaining stock analytic data
- Displaying rolling revenue for a given time period

After completing this module, students will be able to:

- Describe the Cosmos DB change feed, and how to use it to process data efficiently.
- Explain how to incorporate Cosmos DB into streaming solutions such as Apache Kafka, Apache Spark, and Azure Stream Analytics.