DP-200T01: Implementing an Azure Data Solution

Course Details

Course Code: DP-200T01

Duration: 3 days

Notes:

- This course syllabus should be used to determine whether the course is appropriate for the students, based on their current skills and technical training needs.
- Course content, prices, and availability are subject to change without notice.
- Terms and Conditions apply

Elements of this syllabus are subject to change.

About this course

In this course, the students will implement various data platform technologies into solutions that are in line with business and technical requirements including on-premises, cloud, and hybrid data scenarios incorporating both relational and No-SQL data. They will also learn how to process data using a range of technologies and languages for both streaming and batch data.

The students will also explore how to implement data security including authentication, authorization, data policies and standards. They will also define and implement data solution monitoring for both the data storage and data processing activities. Finally, they will manage and troubleshoot Azure data solutions which includes the optimization and disaster recovery of big data, batch processing and streaming data solutions.

Audience Profile

The primary audience for this course is data professionals, data architects, and business intelligence professionals who want to learn about the data platform technologies that exist on Microsoft Azure.

The secondary audience for this course is individuals who develop applications that deliver content from the data platform technologies that exist on Microsoft Azure.

Prerequisites

In addition to their professional experience, students who take this training should have technical knowledge equivalent to the following courses:

- Azure fundamentals

Module 1: Azure for the Data Engineer

This module explores how the world of data has evolved and how cloud data platform technologies are providing new opportunities for business to explore their data in different ways. The student will gain an overview of the various data platform technologies that are available, and how a Data Engineers role and responsibilities has evolved to work in this new world to an organization benefit

Lessons
- Explain the evolving world of data
- Survey the services in the Azure Data Platform
- Identify the tasks that are performed by a Data Engineer
- Describe the use cases for the cloud in a Case Study

Lab : Azure for the Data Engineer
- Identify the evolving world of data
- Determine the Azure Data Platform Services
- Identify tasks to be performed by a Data Engineer
- Finalize the data engineering deliverables

After completing this module, students will be able to:

- Explain the evolving world of data
- Survey the services in the Azure Data Platform
- Identify the tasks that are performed by a Data Engineer
- Describe the use cases for the cloud in a Case Study

Module 2: Working with Data Storage

This module teaches the variety of ways to store data in Azure. The Student will learn the basics of storage management in Azure, how to create a Storage Account, and how to choose the right model for the data you want to store in the cloud. They will also understand how data lake storage can be created to support a wide variety of big data analytics solutions with minimal effort.

Lessons
- Choose a data storage approach in Azure
- Create an Azure Storage Account
- Explain Azure Data Lake storage
- Upload data into Azure Data Lake

Lab : Working with Data Storage
- Choose a data storage approach in Azure
- Create a Storage Account
- Explain Data Lake Storage
- Upload data into Data Lake Store

After completing this module, students will be able to:

- Choose a data storage approach in Azure
- Create an Azure Storage Account
- Explain Azure Data Lake Storage
- Upload data into Azure Data Lake

Module 3: Enabling Team Based Data Science with Azure Databricks

This module introduces students to Azure Databricks and how a Data Engineer works with it to enable an organization to perform Team Data Science projects. They will learn the fundamentals of Azure Databricks and Apache Spark notebooks; how to provision the service and workspaces and learn how to perform data preparation task that can contribute to the data science project.

Lessons
- Explain Azure Databricks and Machine Learning Platforms
- Describe the Team Data Science Process
- Provision Azure Databricks and workspaces
- Perform data preparation tasks

Lab : Enabling Team Based Data Science with Azure Databricks
- Explain Azure Databricks and Machine Learning Platforms
- Describe the Team Data Science Process
- Provision Azure Databricks and Workspaces
- Perform Data Preparation Tasks

After completing this module, students will be able to:

- Explain Azure Databricks
- Describe the Team Data Science Process
- Provision Azure Databricks and workspaces
- Perform data preparation tasks

Module 4: Building Globally Distributed Databases with Cosmos DB

In this module, students will learn how to work with NoSQL data using Azure Cosmos DB. They will learn how to provision the service, and how they can load and interrogate data in the service using Visual Studio Code extensions, and the Azure Cosmos DB .NET Core SDK. They will also learn how to configure the availability options so that users are able to access the data from anywhere in the world.

Lessons
- Create an Azure Cosmos DB database built to scale
- Insert and query data in your Azure Cosmos DB database
- Provision a .NET Core app for Cosmos DB in Visual Studio Code
- Distribute your data globally with Azure Cosmos DB

Lab : Building Globally Distributed Databases with Cosmos DB
- Create an Azure Cosmos DB
- Insert and query data in Azure Cosmos DB
- Build a .Net Core App for Azure Cosmos DB using VS Code
- Distribute data globally with Azure Cosmos DB

After completing this module, students will be able to:

- Create an Azure Cosmos DB database built to scale
- Insert and query data in your Azure Cosmos DB database
- Build a .NET Core app for Azure Cosmos DB in Visual Studio Code
- Distribute your data globally with Azure Cosmos DB

Module 5: Working with Relational Data Stores in the Cloud

In this module, students will explore the Azure relational data platform options including SQL Database and SQL Data Warehouse. The student will be able explain why they would choose one service over another, and how to provision, connect and manage each of the services.

Lessons
- SQL Database and SQL Data Warehouse
- Provision an Azure SQL database to store data
- Provision and load data into Azure SQL Data Warehouse

Lab : Working with Relational Data Stores in the Cloud
- Explain SQL Database and SQL Data Warehouse
- Create an Azure SQL Database to store data
- Provision and load data into Azure SQL Data Warehouse

After completing this module, students will be able to:

- Explain SQL Database and SQL Data Warehouse
- Provision an Azure SQL database to store application data
- Provision and load data in Azure SQL Data Warehouse
- Import data into Azure SQL Data Warehouse using PolyBase

Module 6: Performing Real-Time Analytics with Stream Analytics

In this module, students will learn the concepts of event processing and streaming data and how this applies to Events Hubs and Azure Stream Analytics. The students will then set up a stream analytics job to stream data and learn how to query the incoming data to perform analysis of the data. Finally, you will learn how to manage and monitor running jobs.

Lessons
- Explain data streams and event processing
- Querying streaming data using Stream Analytics
- How to process data with Azure Blob and Stream Analytics
- How to process data with Event Hubs and Stream Analytics

Lab : Performing Real-Time Analytics with Stream Analytics
- Explain data streams and event processing

- Querying streaming data using Stream Analytics
- Process data with Azure Blob and Stream Analytics
- Process data with Event Hubs and Stream Analytics

After completing this module, students will be able to:

- Explain data streams and event processing
- Querying streaming data using Stream Analytics
- How to process data with Event Hubs and Stream Analytics
- How to process data with Azure Blob and Stream Analytics

Module 7: Orchestrating Data Movement with Azure Data Factory

In this module, students will learn how Azure Data factory can be used to orchestrate the data movement and transformation from a wide range of data platform technologies. They will be able to explain the capabilities of the technology and be able to set up an end to end data pipeline that ingests and transforms data.

Lessons
- Explain how Azure Data Factory works
- Create Linked Services and datasets
- Create pipelines and activities
- Azure Data Factory pipeline execution and triggers

Lab : Orchestrating Data Movement with Azure Data Factory
- Explain how Data Factory Works
- Create Linked Services and Datasets
- Create Pipelines and Activities
- Azure Data Factory Pipeline Execution and Triggers

After completing this module, students will be able to:

- Explain how Azure Data Factory works
- Create Linked Services and Datasets
- Create Pipelines and Activities
- Azure Data Factory pipeline execution and triggers

Module 8: Securing Azure Data Platforms

In this module, students will learn how Azure Storage provides a multi-layered security model to protect your data. The students will explore how security can range from setting up secure networks and access keys, to defining permission through to monitoring with Advanced Threat Detection.

Lessons
- Configuring Network Security
- Configuring Authentication
- Configuring Authorization
- Auditing Security

Lab : Securing Azure Data Platforms
- Configure network security
- Configure Authentication
- Configure Authorization
- Explore SQL Server Books Online

After completing this module, students will be able to:

- Configure Authentication
- Use storage account keys
- Use shared access signatures
- Configure Authorization
- Control network access
- Understand transport-level encryption with HTTPS
- Understand Advanced Threat Detection

Module 9: Monitoring and Troubleshooting Data Storage and Processing

In this module, the student will look at the wide range of monitoring capabilities that are available to provide operational support should there be issue with a data platform architecture. They will explore the data engineering troubleshooting approach and be able to apply this to common data storage and data processing issues.

Lessons
- Data Engineering troubleshooting approach
- Azure Monitoring Capabilities
- Troubleshoot common data issues
- Troubleshoot common data processing issues

Lab : Monitoring and Troubleshooting Data Storage and Processing
- Explain the Data Engineering troubleshooting approach
- Explain the monitoring capabilities that are available
- Troubleshoot common data storage issues
- Troubleshoot common data processing issues

After completing this module, students will be able to:

- Explain the monitoring capabilities that are available
- Explain the Data Engineering troubleshooting approach
- Troubleshoot common data storage issues
- Troubleshoot common data processing issues

Module 10: Integrating and Optimizing Data Platforms

In this module, the student will explore the various ways in which data platforms can be integrated based upon different business requirements. They will also explore the various ways in which data platforms can be optimized from a storage and data processing perspective to improve data loads. Finally, disaster recovery options are revealed to ensure business continuity.

Lessons
- Integrating data platforms
- Optimizing data stores
- Optimize streaming data
- Manage disaster recovery

Lab : Integrating and Optimizing Data Platforms
- Integrate Data Platforms
- Optimize Data Stores
- Optimize Streaming Data
- Manage Disaster recovery

After completing this module, students will be able to:

- Integrate data platforms
- Optimize relational data stores
- Optimize NoSQL data stores
- Optimize Streaming data stores

- Manage disaster recovery