1. DP-200 - Implementing an Azure Data Solution

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
- Work with Azure Databricks
- Read data with Azure Databricks
- Perform transformations with Azure Databricks

Lab : Enabling Team Based Data Science with Azure Databricks

- Explain Azure Databricks
- Work with Azure Databricks
- Read data with Azure Databricks
- Perform transformations with Azure Databricks

After completing this module, students will be able to:

- Explain Azure Databricks
- Work with Azure Databricks
- Read data with Azure Databricks
- Perform transformations with Azure Databricks

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
- Build 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

- Use Azure SQL Database
- Describe Azure SQL Data Warehouse
- Creating and Querying an Azure SQL Data Warehouse
- Use PolyBase to Load Data into Azure SQL Data Warehouse

Lab : Working with Relational Data Stores in the Cloud

- Use Azure SQL Database
- Describe Azure SQL Data Warehouse
- Creating and Querying an Azure SQL Data Warehouse
- Use PolyBase to Load Data into Azure SQL Data Warehouse

After completing this module, students will be able to:

- Use Azure SQL Database
- Describe Azure Data Warehouse
- Creating and Querying an Azure SQL Data Warehouse
- Using PolyBase to Load Data into Azure SQL Data Warehouse

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
- Data Ingestion with Event Hubs
- Processing Data with Stream Analytics Jobs

Lab : Performing Real-Time Analytics with Stream Analytics

- Explain data streams and event processing
- Data Ingestion with Event Hubs
- Processing Data with Stream Analytics Jobs

After completing this module, students will be able to:

- Explain data streams and event processing
- Data Ingestion with Event Hubs
- Processing Data with Stream Analytics Jobs

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
- Azure Data Factory Components
- Azure Data Factory and Databricks

Lab : Orchestrating Data Movement with Azure Data Factory

- Explain how Data Factory Works
- Azure Data Factory Components

- Azure Data Factory and Databricks

After completing this module, students will be able to:

- Azure Data Factory and Databricks
- Azure Data Factory Components
- Explain how Azure Data Factory works

Module 8: Securing Azure Data Platforms

In this module, students will learn how Azure 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 across a range of data stores.

Lessons

- An introduction to security
- Key security components
- Securing Storage Accounts and Data Lake Storage
- Securing Data Stores
- Securing Streaming Data

Lab : Securing Azure Data Platforms

- An introduction to security
- Key security components
- Securing Storage Accounts and Data Lake Storage
- Securing Data Stores and Streaming Data

After completing this module, students will be able to:

- An introduction to security
- Key security components
- Securing Storage Accounts and Data Lake Storage
- Securing Data Stores and Streaming Data

Module 9: Monitoring and Troubleshooting Data Storage and Processing

In this module, the student will get an overview of the range of monitoring capabilities that are available to provide operational support should there be issue with a data platform architecture. They will explore the common data storage and data processing issues. Finally, disaster recovery options are revealed to ensure business continuity.

Lessons

- Explain the monitoring capabilities that are available
- Troubleshoot common data storage issues
- Troubleshoot common data processing issues
- Manage disaster recovery

Lab : Monitoring and Troubleshooting Data Storage and Processing

- Explain the monitoring capabilities that are available
- Troubleshoot common data storage issues
- Troubleshoot common data processing issues
- Manage disaster recovery

After completing this module, students will be able to:

- Explain the monitoring capabilities that are available
- Troubleshoot common data storage issues
- Troubleshoot common data processing issues
- Manage disaster recovery

2. DP-201 - Designing an Azure Data Solution

Module 1: Data Platform Architecture Considerations

In this module, the students will learn how to design and build secure, scalable and performant solutions in Azure by examining the core principles found in every good architecture. They will learn how using key principles throughout your architecture regardless of technology choice, can help you design, build, and continuously improve your architecture for an organizations benefit.

Lessons

- Core Principles of Creating Architectures
- Design with Security in Mind
- Performance and Scalability
- Design for availability and recoverability
- Design for efficiency and operations
- Case Study

Lab : Case Study

- Design with security in mind
- Consider performance and scalability
- Design for availability and recoverability
- Design for efficiency and operations

After completing this module, students will be able to:

- Design with Security in mind
- Consider performance and scalability
- Design for availability and recoverability
- Design for efficiency and operations

Module 2: Azure Batch Processing Reference Architectures

In this module, the student will learn the reference design and architecture patterns for dealing with the batch processing of data. The student will be exposed to dealing with the movement of data from on-premises systems into a cloud data warehouse and how it can be automated. The student will also be exposed to an AI architecture and how the data platform can integrate with an AI solution.

Lessons

- Lambda architectures from a Batch Mode Perspective
- Design an Enterprise BI solution in Azure
- Automate enterprise BI solutions in Azure
- Architect an Enterprise-grade Conversational Bot in Azure

Lab : Architect an Enterprise-grade Conversational Bot in Azure

- Designing an Enterprise BI solution in Azure
- Automate an Enterprise BI solution in Azure
- Automate an Enterprise BI solution in Azure

After completing this module, students will be able to:

- Core Principles of Creating Architectures
- Describe Lambda architectures from a Batch Mode Perspective
- Design an Enterprise BI solution in Azure
- Automate enterprise BI solutions in Azure
- Architect an Enterprise-grade conversational bot in Azure

- Case study

Module 3: Azure Real-Time Reference Architectures

In this module, the student will learn the reference design and architecture patterns for dealing with streaming data. They will learn how streaming data can be ingested by Event Hubs and Stream Analytics to deliver real-time analysis of data. They will also explore a data science architecture the streams data into Azure Databricks to perform trend analysis. They will finally learn how an Internet of Things (IoT) architecture will require data platform technologies to store data.

Lessons

- Lambda architectures for a Real-Time Perspective
- Architect a stream processing pipeline with Azure Stream Analytics
- Design a stream processing pipeline with Azure Databricks
- Create an Azure IoT reference architecture

Lab : Azure Real-Time Reference Architectures

- Architect a stream processing pipeline with Azure Stream Analytics
- Design a stream processing pipeline with Azure Databricks
- Create an Azure IoT reference architecture

After completing this module, students will be able to:

- Lambda architectures for a Real-Time Mode Perspective
- Architect a stream processing pipeline with Azure Stream Analytics
- Design a stream processing pipeline with Azure Databricks
- Create an Azure IoT reference architecture

Module 4: Data Platform Security Design Considerations

In this module, the student will learn how to incorporate security into an architecture design and learn the key decision points in Azure provides to help you create a secure environment through all the layers of your architecture.

Lessons

- Defense in Depth Security Approach
- Identity Management
- Infrastructure Protection
- Encryption Usage
- Network Level Protection
- Application Security

Lab : Data Platform Security Design Considerations

- Defense in Depth Security Approach
- Identity Protection

After completing this module, students will be able to:

- Defense in Depth Security Approach
- Identity Management
- Infrastructure Protection
- Encryption Usage
- Network Level Protection
- Application Security

Module 5: Designing for Resiliency and Scale

In this module, student will learn scaling services to handle load. They will learn how identifying network bottlenecks and optimizing your storage performance are important to ensure your users have the best experience. They will also learn how to handle infrastructure and service failure, recover from the loss of data, and recover from a disaster by designing availability and recoverability into your architecture.

Lessons

- Adjust Workload Capacity by Scaling
- Optimize Network Performance
- Design for Optimized Storage and Database Performance
- Identifying Performance Bottlenecks
- Design a Highly Available Solution
- Incorporate Disaster Recovery into Architectures
- Design Backup and Restore strategies

Lab : Designing for Resiliency and Scale

- Adjust Workload Capacity by Scaling
- Design for Optimized Storage and Database Performance
- Design a Highly Available Solution
- Incorporate Disaster Recovery into Architectures

After completing this module, students will be able to:

- Adjust Workload Capacity by Scaling
- Optimize Network Performance
- Design for Optimized Storage and Database Performance
- Identifying Performance Bottlenecks
- Design a Highly Available Solution
- Incorporate Disaster Recovery into Architectures
- Design Backup and Restore strategies

Module 6: Design for Efficiency and Operations

In this module, students will learn how to design an Azure architecture that is operationally-efficient and minimizes costs by reducing spend, they will understand how to design architectures that eliminates waste and gives them full visibility into what is being utilized in your organizations Azure environment.

Lessons

- Maximizing the Efficiency of your Cloud Environment
- Use Monitoring and Analytics to Gain Operational Insights
- Use Automation to Reduce Effort and Error

Lab : Design for Efficiency and Operations

- Maximize the Efficiency of your Cloud Environment
- Use Monitoring and Analytics to Gain Operational Insights
- Use Automation to Reduce Effort and Error

After completing this module, students will be able to:

- Maximize the Efficiency of your Cloud Environment
- Use Monitoring and Analytics to Gain Operational Insights
- Use Automation to Reduce Effort and Error