

DP-300 : Administering Relational Database on Microsoft Azure

Course outline

Module 1: The Role of the Azure Database Administrator

This module explores the role of a database administrator in the world of Azure. It also provides some foundational information relevant to the overall content. This includes a review of the various SQL Server-based options (SQL Server in a VM, Managed Instances, and Azure SQL Database.) Students will learn why compatibility level is a crucial concept when working with SQL databases in Azure. Students are also introduced to other database platforms available on Azure in addition to those based on SQL Server, in particular PostgreSQL and MySQL

Lessons

- Azure Data Platform Roles
- Azure Database Platforms and Options
- SQL Server Compatibility Levels
- Azure Preview Features

Lab : Using the Azure Portal and SQL Server Management Studio

- Provision a SQL Server on an Azure Virtual Machine
- Connect to SQL Server and Restore a Backup

After completing this module, students will:

- Understand the role of Azure Database Administrator as it fits in with other data platform roles
- Be able to describe the key differences between the SQL Server-based database options
- Be able to describe the difference between versions and compatibility levels
- Know how to enable and disable preview features

Module 2: Plan and Implement Data Platform Resources

This module introduces methods for deploying data platform resources in Azure. You will learn about options for both upgrading and migrating existing SQL databases to Azure. You will learn how to set up Azure resources to host SQL Server on a Virtual Machine, a Managed Instance, Azure SQL Database and either PostgreSQL or MySQL. You will learn how to determine which options are best based on specific requirements including the High Availability and Disaster Recovery (HADR) needs. They will learn to calculate resource requirements and create templates for their deployments.

Lessons

- Deploying SQL Server using IaaS
- Deploying SQL Server using PaaS
- Deploying Open Source Database Solutions on Azure

Lab : Deploying Azure SQL Database

- Deploy a VM using an ARM template
- Configure resources needed prior to creating a database
- Deploy an Azure SQL Database
- Register the Azure SQL Database instance in Azure Data Studio and validate connectivity
- Deploy PostgreSQL or MySQL using a client tool to validate connectivity

After completing this module, students will:

- Deploy resources using manual methods
- Recommend an appropriate database offering based on requirements
- Configure database resources
- Evaluate and implement a strategy for moving a database to Azure

Module 3: Implement a Secure Environment

This module explores the practices of securing your SQL Server Database as well as an Azure SQL database. This includes a review of the various SQL Server-based options as well as the various Azure options for securing Azure SQL Database as well as the databases with reside within Azure SQL Database. . Students will learn why security is crucial when working with databases. . Students are also introduced to other database platforms available on Azure in addition to those based on SQL Server, in particular, Azure Database for MariaDB/MySQL/PostgreSQL

Lessons

- Configure Database Authentication
- Configure Database Authorization
- Implement Security for Data at Rest
- Implement Security for Data in Transit
- Implement Compliance Controls for Sensitive Data

Lab : Implement a Secure Environment

- Configure a server-based firewall rule using the Azure Portal
- Authorize Access to Azure SQL Database with Azure Active Directory
- Enable Advanced Data Security and Data Classification
- Manage access to database objects

After completing this module, students will:

- Understand the differences between Windows, SQL Server and Azure Active Directory Authentication
- Describe and configure both data-at-rest encryption solutions as well as data-in-transit encryption
- Implement a data sensitivity solution

Module 4: Monitor and Optimize Operational Resources

This module will teach you about resource optimization for your databases created using either IaaS or PaaS services. The module also covers monitoring server and hardware resources. It will familiarize you with the various tools available for monitoring performance and establishing a baseline. You will learn how to interpret performance metrics for the most critical resources. You will also learn how to troubleshoot database performance using Azure SQL Database Intelligent Insights.

Lessons

- Baselines and Performance Monitoring
- Major Causes of Performance Issues
- Configuring Resources for Optimal Performance
- User Database Configuration
- Performance-related Maintenance Tasks

Lab : Monitor and Optimize Resources

- Isolate CPU Problems
- Use Query Store observe blocking problems
- Detect and correct fragmentation issues

After completing this module, students will:

- Monitor activity and compare to a baseline
- Define maintenance tasks related to performance
- Identify major causes of performance problems
- Configure resources for optimal performance
- Configure a user database for optimal performance

Module 5: Optimize Query Performance

Query execution plans are potentially the most important aspect of database performance. Improving bad plans is certainly an area where a small amount of effort can bring huge improvements. While hardware issues can limit query performance, improving hardware usually yields performance improvements in the 10-20% range, at most. More commonly database administrators encounter queries that are not optimized, have stale or missing statistics, have

missing indexes, or poor database design choices that lead to the database engine doing more work than is necessary to return results for a given query. Improving the plans can sometimes yield performance improvements in the 100-200% range or even more, meaning that after improving a plan with better indexes or statistics, a query could run twice or three times as fast! This module provides details on how to analyze individual query performance and determine where improvements can be made.

Lessons

- Understanding SQL Server Query Plans
- Explore Performance-based Database Design
- Evaluate Performance Improvements

Lab : Query Performance Troubleshooting

- Identify issues with database design AdventureWorks2017
- Isolate problem areas in poorly performing queries in AdventureWorks2017
- Use Query Store to detect and handle regression in AdventureWorks2017
- Use query hints to impact performance in AdventureWorks2017

After completing this module, students will:

- Analyze query plans and identify problem areas
- Evaluate potential query improvements
- Review table and index design
- Determine whether query or design changes have had a positive effect

Module 6: Automation of Tasks

A common goal for database administrators in many environments is to automate as many of their repetitive tasks. This can be as simple as using scripting to automate a backup process, and as complex as building a fully automated alerting system. This module provides details of automating tasks to simplify the DBA's job. Methods include scheduling tasks for regular maintenance jobs, as well as multi-instance administration and configuration of notifications for task success or failure or non-completion.

Lessons

- Setting up Automatic Deployment
- Defining Scheduled Tasks
- Configuring Extended Events
- Managing Azure PaaS resources Using Automated Methods

Lab : Automating Tasks

- Deploy an Azure template from a Quickstart template on GitHub

- Configure notifications based on performance metrics
- Deploy an Azure Automation Runbook (or elastic job) to rebuild indexes on an Azure SQL Database

After completing this module, students will:

- Deploy resources using automated deployment scripts
- Create scheduled tasks
- Create notifications and alerts
- Configure automation for PaaS services

Module 7: Plan and Implement a High Availability and Disaster Recovery Environment

Data must be available when the business needs it. That means the solutions hosting the data must be designed with availability and recoverability in mind. Suppose you work for a company that sells widgets both in stores and online. Your main application uses a highly transactional database for orders. What would happen if the server or platform hosting the transactional database had a problem that made it unavailable or inaccessible for some reason? What impact would it have on the business? If the right solution is put in place, the database would come online in a reasonable timeframe with minimal effort, thus allowing business to continue with little-to-no impact. This module and its associated lab cover configuring, testing, and managing a solution for high availability and disaster recovery (HADR) in Azure, for both Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) deployments. This module will not only cover basic requirements, but also the various options available to achieve HADR.

Lessons

- High Availability and Disaster Recovery Strategies
- IaaS Platform and Database Tools for HADR
- PaaS Platform and Database Tools for HADR
- Database Backup and Recovery

Lab : Plan and Implement a High Availability and Disaster Recovery Environment

- Create an Always On Availability Group
- Enable Geo-Replication for Azure SQL Database
- Backup to URL and Restore from URL

After completing this module, students will understand:

- The difference between recovery time and recovery point objectives
- The available HADR options for both IaaS and PaaS
- The considerations for planning and configuring HADR solutions including how backup and restore fit
- The factors that comprise a HADR strategy

- How to configure a high availability solution via a hands-on lab