

DP-300T00-A: Implement scalable database solutions using Azure SQL

Overview

This course provides students with the knowledge and skills to create a SQL Server database infrastructure for cloud, on-premises and hybrid relational databases and who work with the Microsoft PaaS relational database offerings. Additionally, it will be of use to individuals who develop applications that deliver content from SQL-based relational databases.

Audience Profile

The audience for this course is data professionals managing data and databases who want to learn about administering the data platform technologies that are available on Microsoft Azure. This course is also valuable for data architects and application developers who need to understand what technologies are available for the data platform with Azure and how to work with those technologies through applications.

Duration : 4 Days

Table of Contents :

Plan and implement data platform resources

- Prepare to maintain SQL Server-based databases on Azure
 - Describe Microsoft Intelligent Data Platform roles
 - Understand SQL Server in an Azure virtual machine
 - Design Azure SQL Database for cloud-native applications
 - Explore Azure SQL Managed Instance
- Deploy IaaS solutions with Azure SQL

- Explain IaaS options to deploy SQL Server in Azure
- Understand hybrid scenarios
- Explore performance and security
- Explain high availability and disaster recovery options
- Deploy PaaS solutions with Azure SQL
 - Explain PaaS options for deploying SQL Server in Azure
 - Explore single SQL database
 - Deploy SQL database elastic pool
 - Understand SQL database hyperscale
 - Examine SQL managed instance
- Migrate SQL Server workloads to Azure SQL Database
 - Choose the right Azure SQL Database feature
 - Use Azure SQL migration extension to migrate to Azure SQL Database
 - Migrate to Azure SQL Database using BACPAC
 - Use an online method to migrate to Azure SQL Database
 - Move data to Azure SQL Database
- Migrate SQL Server workloads to Azure SQL Managed Instance
 - Evaluate migration scenarios
 - Use Log Replay Service LRS to migrate
 - Migrate using Managed Instance link
 - Move data to SQL Managed Instance

Implement a secure environment for a database service

- Configure database authentication and authorization
 - Describe authentication and identities
 - Describe Security Principals

- Describe database and object permissions
- Identify authentication and authorization failures
- Protect data in-transit and at rest
 - Explore Transparent Data Encryption
 - Configure server and database firewall rules
 - Explain object encryption and secure enclaves
 - Enable encrypted connections
 - Describe SQL injection
 - Understand Azure Key Vault
- Implement compliance controls for sensitive data
 - Explore data classification
 - Explore server and database audit
 - Implement Dynamic Data Masking
 - Implement Row Level security
 - Understand Microsoft Defender for SQL
 - Explore Ledger
 - Implement Microsoft Purview

Monitor and optimize operational resources in Azure SQL

- Describe performance monitoring
 - Describe performance monitoring tools
 - Describe critical performance metrics
 - Establish baseline metrics
 - Explore extended events
 - Describe database watcher preview
 - Explore Query Performance Insight

- Configure SQL Server resources for optimal performance
 - Explain how to optimize Azure storage for SQL Server virtual machines
 - Describe virtual machine resizing
 - Optimize database storage
 - Control SQL Server resources
- Configure databases for optimal performance
 - Explore database maintenance checks
 - Describe database scoped configuration options
 - Describe automatic tuning
 - Describe intelligent query processing

Optimize query performance in Azure SQL

- Explore query performance optimization
 - Understand query plans
 - Explain estimated and actual query plans
 - Describe dynamic management views and functions
 - Explore Query Store
 - Identify problematic query plans
 - Describe blocking and locking
- Explore performance-based database design
 - Describe normalization
 - Choose appropriate data types
 - Design indexes
- Evaluate performance improvements
 - Describe wait statistics
 - Tune and maintain indexes

- Understand query hints
- Explore performance scenarios

Automate database tasks for Azure SQL

- Automate database deployment
 - Describe deployment models in Azure
 - Automate deployment by using Azure Resource Manager templates and Bicep
 - Automate deployment by using PowerShell
 - Automate deployment by using Azure CLI
- Create and manage SQL Agent jobs
 - Create a SQL Server maintenance plan
 - Describe task status notifications
- Manage Azure PaaS tasks using automation
 - Explore Elastic jobs
 - Understand Azure Automation
 - Build an automation runbook
 - Automate database workflows by using Logic Apps
 - Monitor automated tasks

Plan and implement a high availability and disaster recovery environment

- Describe high availability and disaster recovery strategies
 - Describe recovery time objective and recovery point objective
 - Explore high availability and disaster recovery options
 - Describe Azure high availability and disaster recovery features for Azure Virtual Machines

- Describe high availability and disaster recovery for PaaS deployments
- Explore high availability and disaster recovery solution for IaaS
- Describe hybrid solutions
- Explore IaaS and PaaS solutions for high availability and disaster recovery
 - Describe failover clusters in Windows Server
 - Configure Always-on availability groups
 - Describe active geo-replication for Azure SQL Database
 - Explore auto-failover groups for Azure SQL Database and Azure SQL Managed Instance
 - Monitor availability
- Back up and restore databases
 - Back up and restore SQL Server running on Azure virtual machines
 - Back up a SQL Server virtual machine
 - Back up and restore a database for SQL Database and SQL Managed Instance