

Microsoft Azure Data Fundamentals: DP-900

EXAM DESIGN

Audience Profile

Candidates for this exam should have foundational knowledge of core data concepts and how they are implemented using Microsoft Azure data services.

This exam is intended for candidates beginning to work with data in the cloud.

Candidates should be familiar with the concepts of relational and non-relational data, and different types of data workloads such as transactional or analytical.

Azure Data Fundamentals can be used to prepare for other Azure role-based certifications like Azure Database Administrator Associate or Azure Data Engineer Associate, but it is not a prerequisite for any of them.

Objective Domains

SKILLS MEASURED

- NOTE: The bullets that appear below each of the skills measured are intended to illustrate how we are assessing that skill. This list is not definitive or exhaustive.
- NOTE: In most cases, exams do NOT cover preview features, and some features will only be added to an exam when they are GA (General Availability).

Describe core data concepts (15-20%)

Describe types of core data workloads

- Describe batch data
- Describe streaming data
- Describe the difference between batch and streaming data
- Describe the characteristics of relational data

Describe data analytics core concepts

- Describe data visualization (e.g., visualization, reporting, business intelligence (BI))
- Describe basic chart types such as bar charts and pie charts
- Describe analytics techniques (e.g., descriptive, diagnostic, predictive, prescriptive, cognitive)
- Describe ELT and ETL processing
- Describe the concepts of data processing

Describe how to work with relational data on Azure (25-30%)

Describe relational data workloads

- Identify the right data offering for a relational workload
- Describe relational data structures (e.g., tables, index, views)



Describe relational Azure data services

- Describe and compare PaaS, IaaS, and SaaS solutions
- Describe Azure SQL family of products including Azure SQL Database, Azure SQL Managed Instance, and SQL Server on Azure Virtual Machine
- Describe Azure Synapse Analytics
- Describe Azure Database for PostgreSQL, Azure Database for MariaDB, and Azure Database for MySQL

Identify basic management tasks for relational data

- Describe provisioning and deployment of relational data services
- Describe method for deployment including the Azure portal, Azure Resource Manager templates, Azure PowerShell, and the Azure command-line interface (CLI)
- Identify data security components (e.g., firewall, authentication)
- Identify basic connectivity issues (e.g., accessing from on-premises, access with Azure VNets, access from Internet, authentication, firewalls)
- Identify query tools (e.g., Azure Data Studio, SQL Server Management Studio, sqlcmd utility, etc.)

Describe query techniques for data using SQL language

- Compare Data Definition Language (DDL) versus Data Manipulation Language (DML)
- Query relational data in Azure SQL Database, Azure Database for PostgreSQL, and Azure Database for MySQL

Describe how to work with non-relational data on Azure (25-30%)

Describe non-relational data workloads

- Describe the characteristics of non-relational data
- Describe the types of non-relational data
- Recommend the correct data store
- Determine when to use non-relational data

Describe non-relational data offerings on Azure

- Identify Azure data services for non-relational workloads
- Describe Azure Cosmos DB APIs
- Describe Azure Table storage
- Describe Azure Blob storage
- Describe Azure File storage

Identify basic management tasks for non-relational data

- Describe provisioning and deployment of non-relational data services
- Describe method for deployment including the Azure portal, Azure Resource Manager templates, Azure PowerShell, and the Azure command-line interface (CLI)
- Identify data security components (e.g., firewall, authentication, encryption)
- Identify basic connectivity issues (e.g., accessing from on-premises, access with Azure VNets, access from Internet, authentication, firewalls)
- Identify management tools for non-relational data

Describe an analytics workload on Azure (25-30%)

Describe analytics workloads

- Describe transactional workloads
- Describe the difference between a transactional and an analytics workload
- Describe the difference between batch and real time
- Describe data warehousing workloads
- Determine when a data warehouse solution is needed

Describe the components of a modern data warehouse

- Describe Azure data services for modern data warehousing such as Azure Data Lake, Azure Synapse Analytics, Azure Databricks, Storage Gen2, and Azure HDInsight
- Describe modern data warehousing architecture and workload

Describe data ingestion and processing on Azure

- Describe common practices for data loading
- Describe the components of Azure Data Factory (e.g., pipeline, activities, etc.)
- Describe data processing options (e.g., Azure HDInsight, Azure Databricks, Azure Synapse Analytics, Azure Data Factory)

Describe data visualization in Microsoft Power BI

- Describe the role of paginated reporting
- Describe the role of interactive reports
- Describe the role of dashboards
- Describe the workflow in Power BI