

Microsoft Certified: Azure Data Fundamentals – Skills Measured

This document contains the skills measured on the exams associated with this certification. It does not include any upcoming or recent changes that have been made to those skills. For more information about upcoming or recent changes, see the associated exam details page(s).

NOTE: The bullets that follow each of the skills measured are intended to illustrate how we are assessing that skill. This list is NOT definitive or exhaustive.

NOTE: Most questions cover features that are General Availability (GA). The exam may contain questions on Preview features if those features are commonly used.

Exam DP-900: Microsoft Azure Data Fundamentals

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 Machines
- 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 from 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 from 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 Storage Gen2, Azure Synapse Analytics, Azure Databricks, 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