



---

# MICROSOFT DP-100

---

**Microsoft Designing and Implementing a Data Science Solution on  
Azure Certification Questions & Answers**

---

Exam Summary – Syllabus – Questions

---

**DP-100**

**[Microsoft Certified - Azure Data Scientist Associate](#)**

**40-60 Questions Exam - 700/1000 Cut Score - Duration of 120 minutes**

## Table of Contents:

Know Your DP-100 Certification Well: .....	2
Microsoft DP-100 Designing and Implementing a Data Science Solution on Azure Certification Details:.....	2
DP-100 Syllabus:.....	3
Manage Azure resources for machine learning (25-30%).....	3
Run Experiments and Train Models (20-25%) .....	3
Deploy and operationalize machine learning solutions (35-40%) .....	4
Implement Responsible ML (5-10%).....	5
Microsoft DP-100 Sample Questions: .....	6
Study Guide to Crack Microsoft Designing and Implementing a Data Science Solution on Azure DP-100 Exam: .....	11

## Know Your DP-100 Certification Well:

The DP-100 is best suitable for candidates who want to gain knowledge in the Microsoft Microsoft Azure. Before you start your DP-100 preparation you may struggle to get all the crucial Designing and Implementing a Data Science Solution on Azure materials like DP-100 syllabus, sample questions, study guide.

But don't worry the DP-100 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the DP-100 syllabus?
- How many questions are there in the DP-100 exam?
- Which Practice test would help me to pass the DP-100 exam at the first attempt?

Passing the DP-100 exam makes you Microsoft Certified - Azure Data Scientist Associate. Having the Designing and Implementing a Data Science Solution on Azure certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

## Microsoft DP-100 Designing and Implementing a Data Science Solution on Azure Certification Details:

Exam Name	Microsoft Certified - Azure Data Scientist Associate
Exam Code	DP-100
Exam Price	\$165 (USD)
Duration	120 mins
Number of Questions	40-60
Passing Score	700 / 1000
Books / Training	<a href="#">DP-100T01-A: Designing and Implementing a Data Science Solution on Azure</a>
Schedule Exam	<a href="#">Pearson VUE</a>
Sample Questions	<a href="#">Microsoft Designing and Implementing a Data Science Solution on Azure Sample Questions</a>
Practice Exam	<a href="#">Microsoft DP-100 Certification Practice Exam</a>

## DP-100 Syllabus:

Topic	Details
<b>Manage Azure resources for machine learning (25-30%)</b>	
Create an Azure Machine Learning workspace	<ul style="list-style-type: none"> <li>- create an Azure Machine Learning workspace</li> <li>- configure workspace settings</li> <li>- manage a workspace by using Azure Machine Learning studio</li> </ul>
Manage data in an Azure Machine Learning workspace	<ul style="list-style-type: none"> <li>- select Azure storage resources</li> <li>- register and maintain datastores</li> <li>- create and manage datasets</li> </ul>
Manage compute for experiments in Azure Machine Learning	<ul style="list-style-type: none"> <li>- determine the appropriate compute specifications for a training workload</li> <li>- create compute targets for experiments and training</li> <li>- configure Attached Compute resources including Azure Databricks</li> <li>- monitor compute utilization</li> </ul>
Implement security and access control in Azure Machine Learning	<ul style="list-style-type: none"> <li>- determine access requirements and map requirements to built-in roles</li> <li>- create custom roles</li> <li>- manage role membership</li> <li>- manage credentials by using Azure Key Vault</li> </ul>
Set up an Azure Machine Learning development environment	<ul style="list-style-type: none"> <li>- create compute instances</li> <li>- share compute instances</li> <li>- access Azure Machine Learning workspaces from other development environments</li> </ul>
Set up an Azure Databricks workspace	<ul style="list-style-type: none"> <li>- create an Azure Databricks workspace</li> <li>- create an Azure Databricks cluster</li> <li>- create and run notebooks in Azure Databricks</li> <li>- link an Azure Databricks workspace to an Azure Machine Learning workspace</li> </ul>
<b>Run Experiments and Train Models (20-25%)</b>	
Create models by using the Azure Machine Learning Designer	<ul style="list-style-type: none"> <li>- create a training pipeline by using Azure Machine Learning designer</li> <li>- ingest data in a designer pipeline</li> <li>- use designer modules to define a pipeline data flow</li> <li>- use custom code modules in designer</li> </ul>

Topic	Details
Run model training scripts	<ul style="list-style-type: none"> <li>- create and run an experiment by using the Azure Machine Learning SDK</li> <li>- configure run settings for a script</li> <li>- consume data from a dataset in an experiment by using the Azure Machine Learning SDK</li> <li>- run a training script on Azure Databricks compute</li> <li>- run code to train a model in an Azure Databricks notebook</li> </ul>
Generate metrics from an experiment run	<ul style="list-style-type: none"> <li>- log metrics from an experiment run</li> <li>- retrieve and view experiment outputs</li> <li>- use logs to troubleshoot experiment run errors</li> <li>- use MLflow to track experiments</li> <li>- track experiments running in Azure Databricks</li> </ul>
Use Automated Machine Learning to create optimal models	<ul style="list-style-type: none"> <li>- use the Automated ML interface in Azure Machine Learning studio</li> <li>- use Automated ML from the Azure Machine Learning SDK</li> <li>- select pre-processing options</li> <li>- select the algorithms to be searched</li> <li>- define a primary metric</li> <li>- get data for an Automated ML run</li> <li>- retrieve the best model</li> </ul>
Tune hyperparameters with Azure Machine Learning	<ul style="list-style-type: none"> <li>- select a sampling method</li> <li>- define the search space</li> <li>- define the primary metric</li> <li>- define early termination options</li> <li>- find the model that has optimal hyperparameter values</li> </ul>
<b>Deploy and operationalize machine learning solutions (35-40%)</b>	
Select compute for model deployment	<ul style="list-style-type: none"> <li>- consider security for deployed services</li> <li>- evaluate compute options for deployment</li> </ul>
Deploy a model as a service	<ul style="list-style-type: none"> <li>- configure deployment settings</li> <li>- deploy a registered model</li> <li>- deploy a model trained in Azure Databricks to an Azure Machine Learning endpoint</li> <li>- consume a deployed service</li> <li>- troubleshoot deployment container issues</li> </ul>
Manage models in Azure Machine Learning	<ul style="list-style-type: none"> <li>- register a trained model</li> <li>- monitor model usage</li> <li>- monitor data drift</li> </ul>

Topic	Details
Create an Azure Machine Learning pipeline for batch inferencing	<ul style="list-style-type: none"> <li>- configure a ParallelRunStep</li> <li>- configure compute for a batch inferencing pipeline</li> <li>- publish a batch inferencing pipeline</li> <li>- run a batch inferencing pipeline and obtain outputs</li> <li>- obtain outputs from a ParallelRunStep</li> </ul>
Publish an Azure Machine Learning designer pipeline as a web service	<ul style="list-style-type: none"> <li>- create a target compute resource</li> <li>- configure an Inference pipeline</li> <li>- consume a deployed endpoint</li> </ul>
Implement pipelines by using the Azure Machine Learning SDK	<ul style="list-style-type: none"> <li>- create a pipeline</li> <li>- pass data between steps in a pipeline</li> <li>- run a pipeline</li> <li>- monitor pipeline runs</li> </ul>
Apply ML Ops practices	<ul style="list-style-type: none"> <li>- trigger an Azure Machine Learning pipeline from Azure DevOps</li> <li>- automate model retraining based on new data additions or data changes</li> <li>- refactor notebooks into scripts</li> <li>- implement source control for scripts</li> </ul>
<b>Implement Responsible ML (5-10%)</b>	
Use model explainers to interpret models	<ul style="list-style-type: none"> <li>- select a model interpreter</li> <li>- generate feature importance data</li> </ul>
Describe fairness considerations for models	<ul style="list-style-type: none"> <li>- evaluate model fairness based on prediction disparity</li> <li>- mitigate model unfairness</li> </ul>
Describe privacy considerations for data	<ul style="list-style-type: none"> <li>- describe principles of differential privacy</li> <li>- specify acceptable levels of noise in data and the effects on privacy</li> </ul>

## Microsoft DP-100 Sample Questions:

### Question: 1

You manually define security rules to isolate your Azure Machine Learning training processes within an Azure Virtual Network. You need to ensure that only required communications from the Internet are allowed.

Which three actions should you perform?

Each correct answer presents part of the solution. Choose the correct answers.

- a) Define an inbound security rule.
- b) Ensure the source service tag is set to AzureMachineLearning.
- c) Ensure the source service tag is set to BatchNodeManagement.
- d) Configure a source port range of 29876-29877.
- e) Create an outbound security rule.
- f) Configure a destination port range of 22.

**Answer: a, c, d**

### Question: 2

You plan to provision an Azure Machine Learning Basic edition workspace for a data science project. You need to identify the tasks you will be able to perform in the workspace.

Which three tasks will you be able to perform?

Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- a) Create a Compute Instance and use it to run code in Jupyter notebooks.
- b) Create an Azure Kubernetes Service (AKS) inference cluster.
- c) Use the designer to train a model by dragging and dropping pre-defined modules.
- d) Create a tabular dataset that supports versioning.
- e) Use the Automated Machine Learning user interface to train a model.

**Answer: a, b, c, d**

**Question: 3**

You are building a regression model for estimating the number of calls during an event. You need to determine whether the feature values achieve the conditions to build a Poisson regression model.

Which two conditions must the feature set contain?

Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- a) The label data must be a negative value.
- b) The label data must be whole numbers.
- c) The label data must be non-discrete.
- d) The label data must be a positive value.
- e) The label data can be positive or negative.

**Answer: b, d**

**Question: 4**

You are analyzing a dataset by using Azure Machine Learning Studio. You need to generate a statistical summary that contains the p-value and the unique count for each feature column.

Which two modules can you use?

Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- a) Computer Linear Correlation
- b) Export Count Table
- c) Execute Python Script
- d) Convert to Indicator Values
- e) Summarize Data

**Answer: b, e**



**Question: 5**

You have access to structured data used in a previous machine learning project. The data is not registered in a dataset.

You need to use Azure Machine Learning SDK to use the data directly in a machine learning script. What should you do?

Choose the correct answer

- a) Use Python to create a TabularDataset. Specify the path to the data.
- b) Use the get\_context method from the Run class to load the run from the remote environment.
- c) Use the register method from the Dataset class to create a new dataset version.
- d) Use the from\_files method to create a dataset and register the dataset in your workspace.

**Answer: a**

**Question: 6**

You plan to use a Deep Learning Virtual Machine (DLVM) to train deep learning models using Compute Unified Device Architecture (CUDA) computations.

You need to configure the DLVM to support CUDA. What should you implement?

- a) Solid State Drives (SSD)
- b) Computer Processing Unit (CPU) speed increase by using overclocking
- c) Graphic Processing Unit (GPU)
- d) High Random Access Memory (RAM) configuration
- e) Intel Software Extensions (Intel SGX) technology

**Answer: c**

**Question: 7**

You train a machine learning model. You must deploy the model as a real-time inference service for testing.

The service requires low CPU utilization and less than 48 MB of RAM. The compute target for the deployed service must initialize automatically while minimizing cost and administrative overhead.

Which compute target should you use?

- a) Azure Container Instance (ACI)
- b) attached Azure Databricks cluster
- c) Azure Kubernetes Service (AKS) inference cluster
- d) Azure Machine Learning compute cluster

**Answer: a**

**Question: 8**

You use the Two-Class Neural Network module in Azure Machine Learning Studio to build a binary classification model. You use the Tune Model Hyperparameters module to tune accuracy for the model.

You need to configure the Tune Model Hyperparameters module. Which two values should you use?

Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- a) Number of hidden nodes
- b) Learning Rate
- c) The type of the normalizer
- d) Number of learning iterations
- e) Hidden layer specification

**Answer: b, d**

**Question: 9**

You retrain an existing model. You need to register the new version of a model while keeping the current version of the model in the registry. What should you do?

- a) Register a model with a different name from the existing model and a custom property named version with the value 2.
- b) Register the model with the same name as the existing model.
- c) Save the new model in the default datastore with the same name as the existing model. Do not register the new model.
- d) Delete the existing model and register the new one with the same name.

**Answer: a**

**Question: 10**

You use Azure Machine Learning designer to publish an inference pipeline as a web service. During deployment, you create a service principal (SP) and configure authentication.

You need to use the SP while consuming the endpoint. What should you do?

Choose the correct answer.

- a) Use the client secret to retrieve an authentication token.
- b) Use the regen\_key method to regenerate the primary key.
- c) Use the AciWebservice.deploy\_configuration to set auth\_enabled to True.
- d) Use the get\_keys method to retrieve authentication keys.

**Answer: a**

## Study Guide to Crack Microsoft Designing and Implementing a Data Science Solution on Azure DP-100 Exam:

- Getting details of the DP-100 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the DP-100 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the Microsoft provided training for DP-100 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the DP-100 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on DP-100 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

### Reliable Online Practice Test for DP-100 Certification

Make EduSum.com your best friend during your Designing and Implementing a Data Science Solution on Microsoft Azure exam preparation. We provide authentic practice tests for the DP-100 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual DP-100 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the DP-100 exam.

**Start Online practice of DP-100 Exam by visiting URL**

**<https://www.edusum.com/microsoft/dp-100-designing-and-implementing-data-science-solution-microsoft-azure>**