

Study guide for Exam DP-100: Designing and Implementing a Data Science Solution on Azure

QUICK NAVIGATION

Document purpose

About Exam DP-100: Designing and Implementing a Data Science Solution on Azure

Skills measured

Certification journey

Exam overview

Objective domains

Additional study resources

Document purpose

As an attendee of the **Exam Prep session for Exam DP-100: Designing and Implementing a Data Science Solution on Azure**, you can use this guide as a summary of the topics covered and to explore important links and additional resources. The information and materials found here can help you focus your studies as you prepare for the exam.

About Exam DP-100: Designing and Implementing a Data Science Solution on Azure

[Exam DP-100](#) is required to earn the [Azure Data Scientist Associate certification](#).

This exam measures your ability to apply data science and machine learning to implement and run machine learning workloads on Azure.

1 - Exam Prep session study guide for Exam DP-100

© Microsoft Corporation. All rights reserved.

As a candidate for this exam, you should have knowledge and experience in data science and using Azure Machine Learning and Azure Databricks.

Skills measured

For the full list of the skills that the exam measures, along with the level of experience and expertise that you'll need as an exam candidate, check out the [Exam DP-100 skills outline](#).

Certification journey

For an overview of the journey to Microsoft Certification, including prerequisites (if any) and follow-up resources, explore [The journey to Microsoft Certified: Azure Data Scientist Associate](#).

Exam overview

For information on the exam, including the types of questions you may encounter, read [About Microsoft Certification exams](#).

Objective domains

This section itemizes the topics covered in the Exam Prep session and links to Microsoft documentation so you can review the topics in detail.

- [Set up an Azure Machine Learning workspace \(30–35%\)](#)
- [Run experiments and train models \(25–30%\)](#)
- [Optimize and manage models \(20–25%\)](#)
- [Deploy and consume models \(20–25%\)](#)

Set up an Azure Machine Learning workspace (30–35%)

Create an Azure Machine Learning workspace

- [What is Azure Machine Learning?](#)
- [What is an Azure Machine Learning workspace?](#)
- [Manage Azure Machine Learning workspaces](#)
- [Manage assets and resources](#)
- [Azure Machine Learning studio](#)
- [Machine Learning Operationalization \(MLOps\)](#)
- [Continuous integration and delivery \(CI/CD\)](#)
- [Azure Machine Learning and Azure pipelines](#)
- [Azure Machine Learning and GitHub actions](#)
- [Access Control and Permissions](#)

Manage data objects in an Azure Machine Learning workspace

- [Secure data access in Azure Machine Learning](#)
- [Datastores](#)
- [Datasets](#)
- [Create Azure Machine Learning datasets](#)
- [Working with datasets](#)

Manage experiment compute contexts

- [Compute in Azure Machine Learning](#)
- [Selecting training compute targets](#)
- [Attached compute](#)

Run experiments and train models (25–30%)

Create models by using Azure Machine Learning designer

- [Machine Learning designer](#)
- [Create a real-time inference pipeline](#)

Run training scripts in an Azure Machine Learning workspace

- [Train an Azure Machine Learning model](#)
- [Train a machine learning model using ScriptRunConfig](#)

Generate metrics from an experiment run

- [Tracking capabilities in the Azure Machine Learning SDK](#)
- [Enable logging in Azure Machine Learning designer pipelines](#)
- [Monitor and access Machine Learning run logs and metrics](#)
- [Explore run records in the studio](#)
- [Track experiment runs using MLflow tracking](#)
- [Using MLflow](#)

Automate the model training process

- [What are Azure Machine Learning pipelines?](#)
- [Create pipelines with Azure Machine Learning SDK](#)
- [Moving data between Azure Machine Learning pipeline steps](#)
- [Moving data into and between Azure Machine Learning pipeline steps \(Python\)](#)
- [Run Azure Machine Learning pipelines](#)
- [Event-driven workflows](#)

Optimize and manage models (20–25%)

Use Automated Machine Learning to create optimal models

- [What is Automated Machine Learning?](#)
- [Automated machine learning in Azure Machine Learning Studio](#)
- [Train a machine learning model](#)
- [Create review and deploy automated machine learning models with Azure Machine Learning](#)

Use HyperDrive to tune hyperparameters

- [Tune hyperparameters for your model](#)
- [Overview of sampling methods](#)
- [Overview of defining search spaces](#)
- [Tune Hyperparameters with Hyperdrive](#)
- [Overview of termination options](#)
- [Tune hyperparameters for your model with Azure Machine Learning](#)

Use model explainers to interpret models

- [Overview of model interpretability](#)
- [Data interpretability techniques](#)
- [Generate feature importance](#)
- [Assess and mitigate model fairness](#)
- [Differential privacy](#)
- [Evaluate differential privacy](#)
- [Assess fairness in Machine Learning models](#)

Manage models

- [Steps to register a model \(Azure CLI and Python\)](#)
- [Overview of monitoring data](#)
- [Overview of data drifts](#)
- [Detect data drifts on datasets](#)

Deploy and consume models (20–25%)

Create production compute targets

- [Selecting inference compute targets](#)
- [Configure authentication](#)

Deploy a model as a service

- [Define an entry script](#)
- [Define an inference configuration](#)
- [Consume an Azure Machine Learning model deployed as a web service](#)
- [Troubleshoot model deployment](#)
- [Configuring environment containers](#)

Create a pipeline for batch inferencing

- [Create and run the batch scoring pipeline](#)
- [Tutorial: Build an Azure Machine Learning pipeline for batch scoring](#)

Publish a designer pipeline as a web service

- [Create an inference cluster compute target](#)
- [Create a real-time inference pipeline in Machine Learning designer](#)
- [Deploy the real-time endpoint in Machine Learning designer](#)

Additional study resources

In addition to the documentation listed in the previous sections, we offer several resources to help you prepare for the exam and to stay up to speed and engaged with the Azure community. These resources range from formal training to blogs and even interviews with Microsoft team members.

[Course DP-100T01-A: Designing and Implementing a Data Science Solution on Azure](#)

Take a three-day instructor-led course that covers how to operate machine learning solutions at cloud scale using Azure Machine Learning and that combines lectures with practical, hands-on exercises.

[DP-100 learning paths](#)

Don't miss these free, self-paced online resources to help you gain the skills needed to earn your certification.

[DP-100: Designing and Implementing a Data Science Solution on Azure - Microsoft Official Practice Test](#)

Microsoft Official Practice Tests are self-study tools that prepare candidates for the Microsoft required exams. These practice tests are written by subject matter experts and are designed to ensure that all crucial exam objectives are covered in-depth.

[Azure documentation](#)

Stay informed on the latest products, tools, and features, and get information on pricing, partners, support, solutions, and more.

[Azure Community Support](#)

Ask questions, get answers, and connect with Microsoft engineers and Azure community experts.

[Microsoft Learn Community Blog](#)

Get the latest information about certification tests and exam study groups.

[Channel 9](#)

Explore this community site for customers. It includes video channels, discussions, podcasts, screencasts, and interviews.

[Azure Tuesdays with Corey](#)

Corey Sanders answers your questions about Microsoft Azure.

[Azure Fridays](#)

Scott Hanselman, Partner Program Manager, speaks with Azure engineers as they demo capabilities and share insights.

[Microsoft Azure Blog](#)

Keep current on what's happening in Azure, including what's in preview and what's generally available, along with Azure news, updates, and much more.