

Simulate the Real Exam with Microsoft DP-300 Practice Exams



BTW, DOWNLOAD part of TrainingDump DP-300 dumps from Cloud Storage: <https://drive.google.com/open?id=15whI-SXSQ7n62eM3rxVNvHqV3jzpJaA>

Regarding the process of globalization, every fighter who seeks a better life needs to keep pace with its tendency to meet challenges. DP-300 certification is a stepping stone for you to stand out from the crowd. The DP-300 exam guide function as a time-counter, and you can set fixed time to fulfill your task, so that promote your efficiency in real test. The key strong-point of our DP-300 Test Guide is that we impart more important knowledge with fewer questions and answers, with those easily understandable DP-300 study braindumps, you will find more interests in them and experience an easy learning process.

The DP-300 Certification is ideal for database administrators, database developers, and data engineers who want to demonstrate their expertise in managing and maintaining relational databases on the Azure platform. Administering Relational Databases on Microsoft Azure certification validates the candidate's ability to manage and secure database environments, troubleshoot common issues, and optimize database performance for maximum efficiency. It is also a great way to showcase your skills and expertise to potential employers and advance your career in the field of database management.

[**>> Complete DP-300 Exam Dumps <<**](#)

Downloadable DP-300 PDF & Prep DP-300 Guide

In the world of industry, Microsoft Azure certification is the key to a successful career. If you have achieved credential such as Microsoft then it means a bright future is waiting for you. Avail the opportunity of DP-300 dumps at TrainingDump that helps you in achieving good scores in the exam. Due to these innovative methodologies students get help online. The DP-300 Exam Questions Answers are very effective and greatly helpful in increasing the skills of students. They can easily cover the exam topics with more practice due to the unique set of DP-300 exam dumps. The DP-300 certification learning is getting popular with the passage of time.

Microsoft Administering Relational Databases on Microsoft Azure Sample Questions (Q232-Q237):

NEW QUESTION # 232

You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2. ADF1 contains the following pipelines:

P1:Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account P2:Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2 You need to configure P1 and P2 to maximize parallelism and performance. Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

P1:

Set the Copy method to Bulk insert.
Set the Copy method to PolyBase.
Set the Isolation level to Repeatable read.
Set the Partition option to Dynamic range.

P2:

Set the Copy method to Bulk insert.
Set the Copy method to PolyBase.
Set the Isolation level to Repeatable read.
Set the Partition option to Dynamic range.



Answer:

Explanation:

P1:

Set the Copy method to Bulk insert.
Set the Copy method to PolyBase.
Set the Isolation level to Repeatable read.
Set the Partition option to Dynamic range.

P2:

Set the Copy method to Bulk insert.
Set the Copy method to PolyBase.
Set the Isolation level to Repeatable read.
Set the Partition option to Dynamic range.



Explanation

Graphical user interface, text, chat or text message Description automatically generated

P1:

Set the Copy method to Bulk insert.
Set the Copy method to PolyBase.
Set the Isolation level to Repeatable read.
Set the Partition option to Dynamic range.

P2:

Set the Copy method to Bulk insert.
Set the Copy method to PolyBase.
Set the Isolation level to Repeatable read.
Set the Partition option to Dynamic range.



P1: Set the Partition option to Dynamic Range.

The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel.

P2: Set the Copy method to PolyBase

Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store.

(Polybase supports Azure Blob storage and Azure Data Lake Store by default.) Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse>

<https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse>

NEW QUESTION # 233

You have an Azure subscription.

You plan to deploy an Azure SQL managed instance named Server1.

You need to recommend a scaling solution for Server1. The solution must ensure that Server1 can dynamically scale reads and writes.

What should you include in the recommendation?

- A. a DTU based service tier

- B. A Hyperscale based service tier
- **C. A vCore based service tier**
- D. elastic pools

Answer: C

Explanation:

Dynamically scale database resources with minimal downtime - Azure SQL Database & Azure SQL Managed Instance Azure SQL Managed Instance allows you to scale as well:

Azure SQL Managed Instance uses vCores mode and enables you to define maximum CPU cores and maximum of storage allocated to your instance. All databases within the managed instance will share the resources allocated to the instance.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/scale-resources>

NEW QUESTION # 234

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs. Which deployment option and pricing tier should you configure?

- A. Azure SQL Database managed instance Business Critical
- **B. Azure SQL Database Premium**
- C. Azure SQL Database Standard
- D. Azure SQL Database serverless

Answer: B

Explanation:

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.

Incorrect Answers:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

NEW QUESTION # 235

You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:

- * The first six numerals of each customer's phone number must be masked.
- * The last four digits of each customer's phone number must be visible.
- * Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

Microsoft

Exposed Prefix:

0
1
3
5

Padding String:

x
xxxxxx
xx-xxx
xxx-xxx-
x[3]-x[3]

Exposed Suffix:

0
1
3
5

Answer:

Explanation:

Exposed Prefix:

0
1
3
5

Padding String:

x
xxxxxx
xx-xxx
xxx-xxx-
x[3]-x[3]

Exposed Suffix:

0
1
3
5

Explanation:

Box 1: 0

Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix, [padding],suffix Box 2: xxx-xxx Box 3: 5 Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>

NEW QUESTION # 236

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container. Which resource provider should you enable?

- A. Microsoft.Automation
- B. Microsoft.Sql
- C. Microsoft.EventGrid
- D. Microsoft.EventHub

Answer: C

Explanation:

Explanation

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger>

NEW QUESTION # 237

• • • • •

In order to make the exam easier for every candidate, TrainingDump compiled such a study materials that allows making you test and review history performance, and then you can find your obstacles and overcome them. In addition, once you have used this type of DP-300 Exam Question online for one time, next time you can practice in an offline environment. It must be highest efficiently DP-300 exam tool to help you pass the exam.

Downloadable DP-300 PDF: <https://www.trainingdump.com/Microsoft/DP-300-practice-exam-dumps.html>

What's more, part of that TrainingDump DP-300 dumps now are free: <https://drive.google.com/open?id=15whl-SXSQ7n62eM3rxVNvHqV3jzpJaA>