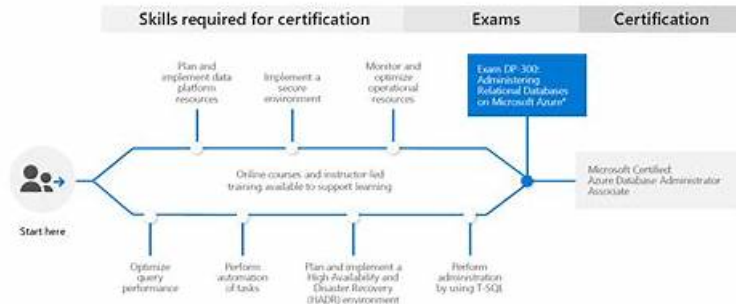


Quiz Microsoft DP-300 - Administering Relational Databases on Microsoft Azure Fantastic Top Exam Dumps

Learning path for Azure Database Administrator Associate



BONUS!!! Download part of PracticeTorrent DP-300 dumps for free: https://drive.google.com/open?id=1RizEKX98If4tMjKZp4Q21Bhu4Bz_O1rp

We offer a money-back guarantee if you fail despite proper preparation and using our product (conditions are mentioned on our guarantee page). This feature gives you the peace of mind to confidently prepare for your Microsoft DP-300 Certification Exam. Our Microsoft DP-300 exam dumps are available for instant download right after purchase, allowing you to start your Microsoft DP-300 preparation immediately.

App online version being suitable to all kinds of digital equipment is supportive to offline exercises on the condition that you practice it without mobile data. These versions of DP-300 test guide make our customers sublimely happy. So they are great DP-300 test guide with high approbation. Our DP-300 Torrent prep is fabulous with inspired points of questions for your reference. After your practice and regular review of our DP-300 exam questions the advancement will be obvious, and your skills of the exam will be improved greatly.

>> DP-300 Top Exam Dumps <<

Efficient and Convenient Preparation with PracticeTorrent's Updated Microsoft DP-300 Exam Questions

PracticeTorrent Microsoft DP-300 practice exam software went through real-world testing with feedback from more than 90,000 global professionals before reaching its latest form. The Microsoft DP-300 Exam Dumps are similar to real exam questions. Our Microsoft DP-300 practice test software is suitable for computer users with a Windows operating system.

Conclusion

The Microsoft DP-300 Exam is becoming an asset for an organization as of now. Using an extensive exam structure, it hones the skills of the test-taker at multiple levels. While you dream of success in DP-300, don't forget to take the help of quality books from Amazon and the official learning paths that we just recommended. They will imbue the best-of-class cognizance about exam topics and help you ace the actual test with confidence.

Microsoft Administering Relational Databases on Microsoft Azure Sample Questions (Q231-Q236):

NEW QUESTION # 231

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

This is not about an external table.

Instead, in an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Note: You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the Get Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity>

NEW QUESTION # 232

You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order.

You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders. What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Create:

- A materialized view in DB1
- A security policy in the Orders table**
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule**
- A table-valued function
- The CONTAINS predicate

practicetorrent.com Microsoft

Answer:

Explanation:



Create:

- A materialized view in DB1
- A security policy in the Orders table
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule
- A table-valued function
- The CONTAINS predicate

Explanation

Create:

- A materialized view in DB1
- A security policy in the Orders table
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule
- A table-valued function
- The CONTAINS predicate

NEW QUESTION # 233

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.

Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted.

You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

| Values | Answer Area |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CustomerKey | <pre> CREATE TABLE [dbo].[FactSales] ([ProductKey] int NOT NULL , [OrderDateKey] int NOT NULL , [CustomerKey] int NOT NULL , [SalesOrderNumber] nvarchar (20) NOT NULL , [OrderQuantity] smallint NOT NULL , [UnitPrice] money NOT NULL) WITH (CLUSTERED COLUMNSTORE INDEX , DISTRIBUTION = [] ([ProductKey]) , PARTITION ([]) RANGE RIGHT FOR VALUES (20170101, 20180101, 20190101, 20200101, 20210101)) </pre> |
| HASH | |
| ROUND_ROBIN | |
| REPLICATE | |
| OrderDateKey | |
| SalesOrderNumber | |

Answer:

Explanation:

Values

Answer Area

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = [HASH] ([ProductKey])
, PARTITION ( [ OrderDateKey ] RANGE RIGHT FOR VALUES
(20170101, 20180101, 20190101, 20200101, 20210101)
)
)
```

Explanation:

Graphical user interface, text, application Description automatically generated

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = [HASH] ([ProductKey])
, PARTITION ( [ OrderDateKey ] RANGE RIGHT FOR VALUES
(20170101, 20180101, 20190101, 20200101, 20210101)
)
)
```

Box 1: HASH

Box 2: OrderDateKey

In most cases, table partitions are created on a date column.

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management.

For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool>

NEW QUESTION # 234

SIMULATION

You need to configure a disaster recovery solution for db1. When a failover occurs, the connection strings to the database must remain the same. The secondary server must be in the West US Azure region.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Explanation:

Configure active geo-replication and failover (Azure SQL Database)

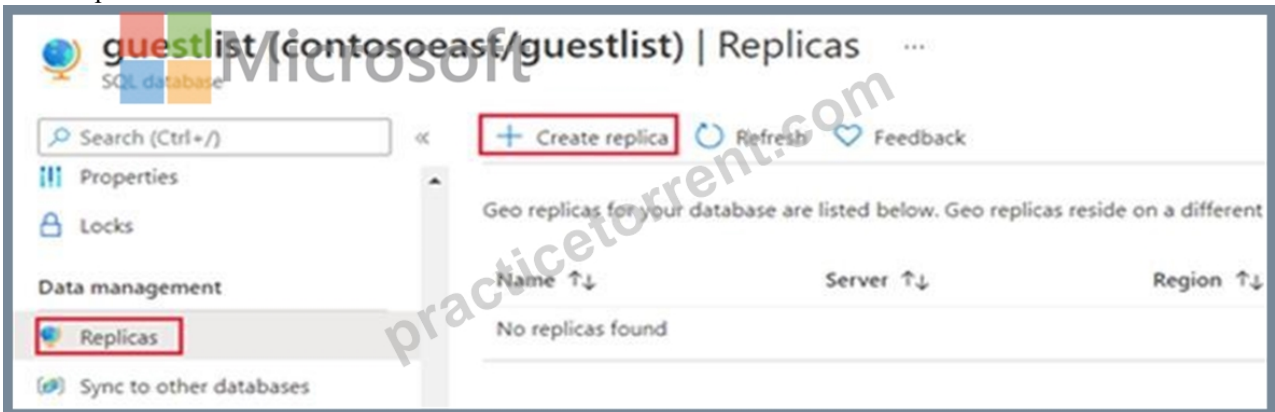
You can configure active geo-replication for Azure SQL Database using the Azure portal or Azure CLI and to initiate failover.

Add a secondary database

The following steps create a new secondary database in a geo-replication partnership.

Step 1: In the Azure portal, browse to the database [Here db1] that you want to set up for geo- replication.

Step 2: On the SQL Database page, select your database [Here db1], scroll to Data management, select Replicas, and then select Create replica.



Step 3: Select or create [Here create] the server for the secondary database, and configure the Compute + storage options if necessary. You can select any region for your secondary server, but we recommend the paired region.

For region select: West US 3

Create SQL Database - Geo Replica

Microsoft

Primary database details

Additional settings will be defaulted where possible based on the the primary database.

Primary database

guestlist

Region

eastus

Database details

Enter required settings for this database, including picking a logical server and configuring the compute and storage resources

Database name

guestlist



Server

contosowest (West US)

Create new

Region

West US

Want to use SQL elastic pool? *

Yes No

Compute + storage *

General Purpose

Standard-series (Gen5), 2 vCores, 32 GB storage, zone redundant disabled

Configure database

Step 4: Click Review + create, review the information, and then click Create.

The secondary database is created and the deployment process begins.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-configure-portal>

NEW QUESTION # 235

You have an Azure SQL managed instance.

