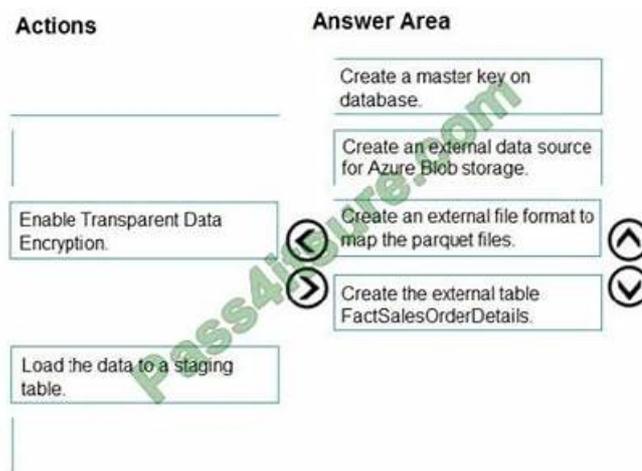


DP-300 Actual Questions, Latest DP-300 Test Format



2026 Latest itPass4sure DP-300 PDF Dumps and DP-300 Exam Engine Free Share: https://drive.google.com/open?id=1hJdz_LgsV2sjm7Rngl6hLbcTTOjgu3fP

Your selection on the right tool to help you pass the DP-300 exam and get the according certification matters a lot for the right DP-300 exam braindumps will spread you a lot of time and efforts. Our DP-300 Study Guide is the most reliable and popular exam product in the market for we only sell the latest DP-300 practice engine to our clients and you can have a free trial before your purchase.

Get information about the some steps to apply for Microsoft DP-300 Exam

In order to apply for the Microsoft DP-300, You have to follow these steps

Step 1: Visit to Microsoft Exam Registration

Step 2: Signup/Login to MICROSOFT account

Step 3: Search for Microsoft DP-300 Certifications Exam

Step 4: Select Date and Center of examination and confirm with payment value of \$165

>> DP-300 Actual Questions <<

Quiz Microsoft - DP-300 - Professional Administering Relational Databases on Microsoft Azure Actual Questions

Our products are the accumulation of professional knowledge worthy practicing and remembering. There are so many specialists who join together and contribute to the success of our DP-300 guide quiz just for your needs. Our responsible and patient staff who has been trained strictly before get down to business and interact with customers. Once you have practiced and experienced the quality of our DP-300 Exam Preparation, you will remember the serviceability and usefulness of them. It explains why our DP-300 practice materials helped over 98 percent of exam candidates get the certificate you dream of successfully. Believe me you can get it too.

Microsoft Administering Relational Databases on Microsoft Azure Sample Questions (Q137-Q142):

NEW QUESTION # 137

Hotspot Question

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales.

You need to initiate a backup of Sales.

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Explanation:

Box 1: TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak' Native database backup in Azure SQL Managed Instance.

You can backup any database using standard BACKUP T-SQL command:

BACKUP DATABASE tpcc2501

TO URL = 'https://myacc.blob.core.windows.net/testcontainer/tpcc2501.bak' WITH COPY_ONLY Box 2: WITH COPY_ONLY SQLMI has "managed" everything, includes backup chain, so user can not break it's chain by using different method than COPY_ONLY.

Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-instance/ba-p/386154>

NEW QUESTION # 138

You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1.

Several users execute adhoc queries to DW1 concurrently.

You regularly perform automated data loads to DW1.

You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run.

What should you do?

- A. Hash distribute the large fact tables in DW1 before performing the automated data loads.
- **B. Assign a larger resource class to the automated data load queries.**
- C. Assign a smaller resource class to the automated data load queries.
- D. Create sampled statistics to every column in each table of DW1.

Answer: B

Explanation:

The performance capacity of a query is determined by the user's resource class.

Smaller resource classes reduce the maximum memory per query, but increase concurrency.

Larger resource classes increase the maximum memory per query, but reduce concurrency.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workloadmanagement>

NEW QUESTION # 139

You have 20 Azure SQL databases provisioned by using the vCore purchasing model.

You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- **A. number of concurrently peaking databases * peak CPU utilization per database**
- B. maximum number of concurrent sessions for all the databases
- **C. total size of all the databases**
- D. geo-replication support
- **E. total number of databases * average CPU utilization per database**

Answer: A,C,E

Explanation:

Section: [none]

Explanation:

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model: $\text{MAX}(\langle \text{Total number of DBs} \times \text{average vCore utilization per DB} \rangle,$

$\text{A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.}$

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview>

NEW QUESTION # 140

You need to identify the cause of the performance issues on SalesSQLDb1.

Which two dynamic management views should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. [sys.dm_pdw_nodes_os_wait_stats](#)
- B. [sys.dm_tran_locks](#)
- C. [sys.dm_exec_compute_node_errors](#)
- D. [sys.dm_exec_requests](#)
- E. [sys.dm_cdc_errors](#)
- F. [sys.dm_pdw_nodes_tran_locks](#)

Answer: A,F

Explanation:

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

A: Use [sys.dm_pdw_nodes_tran_locks](#) instead of [sys.dm_tran_locks](#) from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information.

```
SELECT
t1.resource_type,
t1.resource_database_id,
t1.resource_associated_entity_id,
t1.request_mode,
t1.request_session_id,
t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
ON t1.lock_owner_address = t2.resource_address;
```

Note: Depending on the system you're working with you can access these wait statistics from one of three locations:

[sys.dm_os_wait_stats](#): for SQL Server

[sys.dm_db_wait_stats](#): for Azure SQL Database

[sys.dm_pdw_nodes_os_wait_stats](#): for Azure SQL Data Warehouse

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-locks-transact-sql>

NEW QUESTION # 141

From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays.

The data contains the following columns:

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.

To which table should you add each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/power-bi/guidance/star-schema>

NEW QUESTION # 142

