


DP-300 Advanced Testing Engine | New DP-300 Exam Question



BONUS!!! Download part of ITExamSimulator DP-300 dumps for free: https://drive.google.com/open?id=1tRsE7DHk-lk1_2tF2tSr0AmWJAP8j5v

It doesn't matter if it is the first time you participate in the online training or if you prepare this exam for some time. It is a simple and smart way to prepare the DP-300 practice exam with our latest learning materials. There are free demo and valid questions and answers in our DP-300 Pass Guide. If you spend some time and pay attention to DP-300 test answers, there is no reason to not pass test and get the certification.

Microsoft DP-300 Exam Syllabus Topics:

Section	Objectives
Perform automation	- Automate database tasks • 1. Implement scheduled maintenance • 2. Use Azure Automation
Implement a secure environment	- Implement data protection • 1. Encryption at rest and in transit • 2. Data classification and auditing - Manage database authentication and authorization • 1. Configure Azure Active Directory authentication • 2. Manage database roles and permissions
Monitor and optimize operational resources	- Monitor database performance • 1. Identify performance bottlenecks • 2. Use Azure Monitor and Log Analytics - Optimize resources • 1. Scale compute and storage • 2. Tune performance configuration

Plan and implement data platform resources	- Plan and configure security for the data platform • 1. Configure network security • 2. Implement authentication and authorization - Deploy and configure infrastructure • 1. Configure compute and storage resources • 2. Select appropriate Azure database services
Plan and implement a high availability and disaster recovery environment	- Design HA/DR solutions • 1. Implement geo-replication • 2. Configure failover groups - Backup and restore • 1. Automated backups • 2. Point-in-time restore
Optimize query performance	- Analyze query performance • 1. Identify inefficient queries • 2. Use Query Store and execution plans - Improve performance • 1. Query tuning techniques • 2. Index optimization

>> DP-300 Advanced Testing Engine <<

New DP-300 Exam Question, New DP-300 Exam Testking

The DP-300 practice exam we offered is designed with the real questions that will help you in enhancing your knowledge about the DP-300 certification exam. Our online test engine will improve your ability to solve the difficulty of DP-300 Real Questions and get used to the atmosphere of the formal test. Our experts created the valid DP-300 study guide for most of candidates to help them get good result with less time and money.

Microsoft Administering Relational Databases on Microsoft Azure Sample Questions (Q241-Q246):

NEW QUESTION # 241

Drag and Drop Question

You have an Azure virtual machine named Server1 that contains an instance of Microsoft SQL Server 2022 named SQL1. SQL1 contains two databases named DB1 and DB2.

You need to take a snapshot backup of DB1 and DB2. The backup must NOT disrupt the backup chain and must NOT affect other databases.

Which three Transact-SQL commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Commands

```
BACKUP GROUP db1, db2  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

```
BACKUP SERVER  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON  
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON  
(GROUP = (db1, db2));
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

Answer Area

Answer:

Explanation:

Commands

```
BACKUP SERVER  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON  
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

Answer Area

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON  
(GROUP = (db1, db2));
```

```
BACKUP GROUP db1, db2  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF
```

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1.

You need to ensure that app1 can access db1. The solution must use the principle of least privilege.

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 Area

CREATE USER [app1] FROM EXTERNAL PROVIDER

FOR LOGIN app1
FROM EXTERNAL PROVIDER
FROM LOGIN app1
WITHOUT LOGIN

Answer:

Explanation:

CREATE USER [app1] FROM EXTERNAL PROVIDER

FOR LOGIN app1
FROM EXTERNAL PROVIDER
FROM LOGIN app1
WITHOUT LOGIN

Explanation:

CREATE USER [app1] FROM EXTERNAL PROVIDER

FOR LOGIN app1
FROM EXTERNAL PROVIDER
FROM LOGIN app1
WITHOUT LOGIN

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql>

NEW QUESTION # 243

You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.

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.

Table type to store retail store data:

Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Answer:

Explanation:

Table type to store retail store data:


Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Explanation:

Graphical user interface, text, application Description automatically generated



Box 1: Hash

Scenario:

Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.

A hash distributed table can deliver the highest query performance for joins and aggregations on large tables.

Box 2: Round-robin

Scenario:

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads.

These are some scenarios where you should choose Round robin distribution:

- * When you cannot identify a single key to distribute your data.
- * If your data doesn't frequently join with data from other tables.
- * When there are no obvious keys to join.

Reference:

<https://rajanieshkaushikk.com/2020/09/09/how-to-choose-right-data-distribution-strategy-for-azure-synapse/>

NEW QUESTION # 244

You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.

During peak usage, the database will require the following:

- * 24 cores
- * 500 GB of storage
- * 124 GB of memory
- * More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard.

Which service tier should you use during peak usage?

- A. Business Critical
- B. Hyperscale
- C. Premium

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vc-core-single-databases#business-critical>

NEW QUESTION # 245

You have an Azure SQL database named DB1.

You have a table name Table1 that has 20 columns of type CHAR(400). Row compression for Table1 is enabled.

During a database audit, you discover that none of the fields contain more than 150 characters.

You need to ensure that you can apply page compression to Table1.

What should you do?

- A. Configure the columns as sparse.
- B. Change the column type to varchar (MAX).
- C. Change the column type to varchar (200).

- D. Change the column type to nvarchar (MAX).

Answer: C

Explanation:

We reduce the max length of the column from 400 to 200.

Incorrect:

Not A: Sparse column is useful when there are many null columns.

The SQL Server Database Engine uses the SPARSE keyword in a column definition to optimize the storage of values in that column. Therefore, when the column value is NULL for any row in the table, the values require no storage.

Not B, Not C: SQL Server 2005 got around the limitation of 8KB storage size and provided a workaround with varchar(max). It is a non-Unicode large variable-length character data type and can store a maximum of 2

BTW, DOWNLOAD part of ITExamSimulator DP-300 dumps from Cloud Storage: https://drive.google.com/open?id=1tRsE7DHk-llc1_2tF2tSr0AmWJAP8j5v