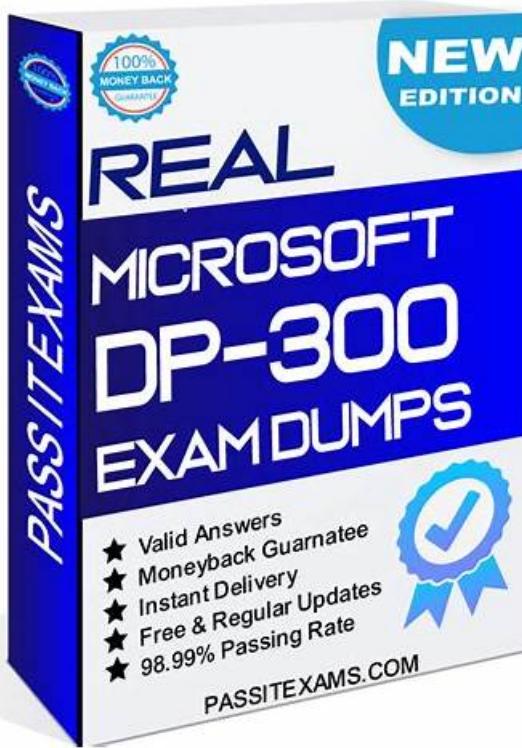


# DP-300 Latest Real Exam & Updated DP-300 Valid Exam Camp Supply you the Best Materials for Administering Relational Databases on Microsoft Azure



What's more, part of that ITEXamSimulator DP-300 dumps now are free: [https://drive.google.com/open?id=1T96-wqcWGqYjhEV6\\_1zdS45qp8x35P](https://drive.google.com/open?id=1T96-wqcWGqYjhEV6_1zdS45qp8x35P)

Our DP-300 study materials are constantly improving themselves. We keep updating them to be the latest and accurate. And we apply the latest technologies to let them applied to the electronic devices. If you have any good ideas, our DP-300 Exam Questions are very happy to accept them. DP-300 learning braindumps are looking forward to having more partners to join this family. We will progress together and become better ourselves.

To become certified in Microsoft DP-300, candidates must pass the certification exam. DP-300 exam consists of multiple-choice questions and performance-based scenarios. Candidates are required to demonstrate their ability to perform tasks related to managing and administering databases on Microsoft Azure. DP-300 exam is available in multiple languages and can be taken online or in-person at a testing center.

Microsoft DP-300 certification exam, also known as Administering Relational Databases on Microsoft Azure, is designed for database administrators who want to showcase their skills and knowledge in managing and maintaining relational databases on the Azure cloud platform. Administering Relational Databases on Microsoft Azure certification exam validates the candidate's ability to design, implement, and manage Azure SQL database solutions, including security, backup, and recovery strategies. DP-300 Exam focuses on various topics such as Azure SQL database creation, monitoring, optimization, and troubleshooting techniques.

>> DP-300 Latest Real Exam <<

## 2026 DP-300: Marvelous Administering Relational Databases on Microsoft Azure Latest Real Exam

Our DP-300 practice engine is admired by all our customers for our experts' familiarity and dedication with the industry all these years. By their help, you can qualify yourself with high-quality DP-300 exam materials. Our experts pass onto the exam candidate their know-how of coping with the exam by our DP-300 Training Questions. And i can say that our DP-300 study guide is the

unique on the market for its high-effective.

## Microsoft DP-300 Exam Syllabus Topics:

| Topic   | Details   |
|---|---|
| <b>Plan and Implement Data Platform Resources (15-20%)</b>                |   |
| Deploy resources by using manual methods                                  | <ul style="list-style-type: none"> <li>-deploy database offerings on selected platforms</li> <li>-configure customized deployment templates</li> <li>-apply patches and updates for hybrid and IaaS deployment</li> </ul>   |
| Recommend an appropriate database offering based on specific requirements | <ul style="list-style-type: none"> <li>-evaluate requirements for the deployment</li> <li>-evaluate the functional benefits/impact of possible database offerings</li> <li>-evaluate the scalability of the possible database offering</li> <li>- evaluate the HA/DR of the possible database offering</li> <li>-evaluate the security aspects of the possible database offering</li> </ul>                       |
| Configure resources for scale and performance                             | <ul style="list-style-type: none"> <li>-configure Azure SQL Database for scale and performance</li> <li>-configure Azure SQL managed instances for scale and performance</li> <li>-configure SQL Server in Azure VMs for scale and performance</li> <li>-calculate resource requirements</li> <li>-evaluate database partitioning techniques, such as database sharding</li> <li>-set up SQL Data Sync</li> </ul> |
| Evaluate a strategy for moving to Azure                                   | <ul style="list-style-type: none"> <li>-evaluate requirements for the migration</li> <li>-evaluate offline or online migration strategies</li> <li>-evaluate requirements for the upgrade</li> <li>-evaluate offline or online upgrade strategies</li> </ul>  |
| Implement a migration or upgrade strategy for moving to Azure             | <ul style="list-style-type: none"> <li>-implement an online migration strategy</li> <li>-implement an offline migration strategy</li> <li>-implement an online upgrade strategy</li> <li>-implement an offline upgrade strategy</li> </ul>  |
| <b>Implement a Secure Environment (15-20%)</b>                            |   |
| Configure database authentication by using platform and database tools    | <ul style="list-style-type: none"> <li>-configure Azure AD authentication</li> <li>-create users from Azure AD identities</li> <li>-configure security principals</li> </ul>  |
| Configure database authorization by using platform and database tools     | <ul style="list-style-type: none"> <li>-configure database and object-level permissions using graphical tools</li> <li>-apply principle of least privilege for all securables</li> </ul>  |
| Implement security for data at rest                                       | <ul style="list-style-type: none"> <li>-implement Transparent Data Encryption (TDE)</li> <li>- implement object-level encryption</li> <li>-implement Dynamic Data Masking</li> <li>-implement Azure Key Vault and disk encryption for Azure VMs</li> </ul>  |
| Implement security for data in transit                                    | <ul style="list-style-type: none"> <li>-configure server and database-level firewall rules</li> <li>-implement Always Encrypted</li> </ul>  |
| Implement compliance controls for sensitive data                          | <ul style="list-style-type: none"> <li>-apply a data classification strategy</li> <li>-configure server and database audits</li> <li>-implement data change tracking</li> <li>-perform a vulnerability assessment</li> </ul>  |
| <b>Monitor and Optimize Operational Resources (15-20%)</b>                |   |
| Monitor activity and performance  | <ul style="list-style-type: none"> <li>-prepare an operational performance baseline</li> <li>-determine sources for performance metrics</li> <li>-interpret performance metrics</li> <li>- configure and monitor activity and performance at the infrastructure, server, service, and database levels</li> </ul>  |

|   |  |
|---|--|
| Implement performance-related maintenance tasks   | <ul style="list-style-type: none"> <li>-implement index maintenance tasks</li> <li>-implement statistics maintenance tasks</li> <li>-configure database auto-tuning</li> <li>-manage storage capacity</li> </ul>   |
| Identify performance-related issues               | <ul style="list-style-type: none"> <li>- configure Query Store to collect performance data</li> <li>-identify sessions that cause blocking</li> <li>-assess growth/fragmentation of databases and logs</li> <li>-assess performance-related database configuration parameters</li> </ul> |
| Configure resources for optimal performance       | <ul style="list-style-type: none"> <li>-configure storage and infrastructure resources</li> <li>-configure server and service account settings for performance</li> <li>-configure Resource Governor for performance</li> </ul>  |
| Configure a user database for optimal performance | <ul style="list-style-type: none"> <li>-implement database-scoped configuration</li> <li>-configure compute resources for scaling</li> <li>-configure Intelligent Query Processing (IQP)</li> </ul>  |

## Optimize Query Performance (5-10%)

### Microsoft Administering Relational Databases on Microsoft Azure Sample Questions (Q187-Q192):

#### NEW QUESTION # 187

Task 12

You need to configure high availability for dbl. The solution must tolerate the loss of an Azure datacenter without data loss or the need to modify application connection strings.

**Answer:**

Explanation:

See the explanation part for the complete Solution

Explanation:

To configure high availability for dbl, you can use the failover groups feature of Azure SQL Database. Failover groups allow you to manage the replication and failover of a group of databases across different regions with the same connection strings1. You can choose all, or a subset of, user databases in a logical server to be replicated to another logical server in a different region. You can also specify the failover policy, such as manual or automatic, and the grace period for data loss.

Here are the steps to create a failover group for dbl:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts dbl.

Select Failover groups in the left menu and click on Add group.

Enter a name for the failover group and select a secondary region that is different from the primary region.

Click on Create a new server and enter the details for the secondary server, such as server name, admin login, password, and subscription.

Click on Select existing database(s) and choose dbl from the list of databases on the primary server.

Click on Configure failover policy and select the failover mode, grace period, and read-write failover endpoint mode according to your preferences.

Click on Create to create the failover group and start the replication of dbl to the secondary server.

Using PowerShell commands:

Install the Azure PowerShell module and log in with your Azure account.

Run the following command to create a new server in the secondary region: `New-AzSqlServer -ResourceGroupName <your-resource-group-name> -ServerName <your-secondary-server-name> -Location "<secondary-region-name>" -SqlAdministratorCredentials $(New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList "<your-admin-login>", $(ConvertTo-SecureString -String "<your-password>" -AsPlainText -Force))` Run the following command to create a new failover group with dbl: `New-AzSqlDatabaseFailoverGroup -ResourceGroupName <your-resource-group-name> -ServerName <your-primary-server-name> -PartnerResourceGroupName <your-resource-group-name> -PartnerServerName <your-secondary-server-name> -FailoverGroupName <your-failover-group-name> -Database dbl -FailoverPolicy <manual-or-automatic> -GracePeriodWithDataLossHours <grace-period-in-hours> -ReadWriteFailoverEndpoint "<enabled-or-disabled>"`

You can modify the parameters of the command according to your preferences, such as the failover policy, grace period, and read-write failover endpoint mode.

These are the steps to create a failover group for dbl

### NEW QUESTION # 188

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 Hyperscale based service tier
- B. a DTU 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:

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 # 189

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases.

The solution must meet the technical requirements and the business requirements.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

The screenshot shows the Microsoft Azure Exam Simulator interface. At the top, there is a Microsoft logo and the text "Purchasing model:" followed by a dropdown menu. The dropdown menu contains three options: "Azure virtual machine reserved instances", "DTU", and "vCore". Below this, there is a "Deployment option:" label followed by another dropdown menu. This second dropdown menu contains three options: "An Azure SQL Database elastic pool", "An Azure SQL Database managed instance", and "A SQL Server Always On availability group".

**Answer:**

Explanation:

Purchasing model:

Microsoft

|  |
|--|
| Azure virtual machine reserved instances |
| DTU                                      |
| vCore                                    |

Deployment option:

|   |
|---|
| An Azure SQL Database elastic pool        |
| An Azure SQL Database managed instance    |
| A SQL Server Always On availability group |

Explanation:

Purchasing model:

|  |
|--|
| Azure virtual machine reserved instances |
| DTU                                      |
| vCore                                    |

Deployment option:

|   |
|---|
| An Azure SQL Database elastic pool        |
| An Azure SQL Database managed instance    |
| A SQL Server Always On availability group |

Box 1: DTU

Scenario:

- \* The 30 new databases must scale automatically.
- \* Once all requirements are met, minimize costs whenever possible.

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore- based purchasing model.

In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.

Box 2: An Azure SQL database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

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

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

#### NEW QUESTION # 190

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.

Exposed Prefix:

|   |
|---|
| 0 |
| 1 |
| 3 |
| 5 |

Padding String:

|           |
|-----------|
| x         |
| xxxxxx    |
| xxx-xxx   |
| xxx-xxx-  |
| x[3]-x[3] |

Exposed Suffix:

|   |
|---|
| 0 |
| 1 |
| 3 |
| 5 |

**Answer:**

Explanation:

Microsoft itexamsimulator.com

Exposed Prefix:

|   |
|---|
| 0 |
| 1 |
| 3 |
| 5 |

Padding String:

|           |
|-----------|
| x         |
| XXXXXX    |
| XXX-XXX   |
| XXX-XXX-  |
| x[3]-x[3] |

Exposed Suffix:

|   |
|---|
| 0 |
| 1 |
| 3 |
| 5 |

Reference:

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

### NEW QUESTION # 191

You are building an Azure Stream Analytics job to retrieve game data.

You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.

How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Microsoft itexamsimulator.com

```

SELECT
    Collect(Score) as HighestScore
    CollectTop(1)OVER(ORDER BY Score Desc)
    Game, MAX(Score)
    TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)
FROM input TIMESTAMP BY CreatedAt
GROUP BY
    Game
    Hopping(minute, 5)
    Tumbling(minute, 5)
    Windows(TumblingWindow(minute, 5), Hopping(minute, 5))

```

Answer:

Explanation:

SELECT

|  |                 |
|--|-----------------|
| Collect(Score)                                       | as HighestScore |
| CollectTop(1)OVER(ORDER BY Score Desc)               |                 |
| Game, MAX(Score)                                     |                 |
| TopOne() OVER(PARTITION BY Game ORDER BY Score Desc) |                 |

FROM input TIMESTAMP BY CreatedAt

GROUP BY

|  |  |
|--|--|
| Game   |  |
| Hopping(minute, 5)                                     |  |
| Tumbling(minute, 5)                                    |  |
| Windows(TumblingWindow(minute, 5), Hopping(minute, 5)) |  |

### Explanation

Graphical user interface, text, application, email Description automatically generated

SELECT

|  |                 |
|--|-----------------|
| Collect(Score)                                       | as HighestScore |
| CollectTop(1)OVER(ORDER BY Score Desc)               |                 |
| Game, MAX(Score)                                     |                 |
| TopOne() OVER(PARTITION BY Game ORDER BY Score Desc) |                 |

FROM input TIMESTAMP BY CreatedAt

GROUP BY

|  |  |
|--|--|
| Game   |  |
| Hopping(minute, 5)                                     |  |
| Tumbling(minute, 5)                                    |  |
| Windows(TumblingWindow(minute, 5), Hopping(minute, 5)) |  |

Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

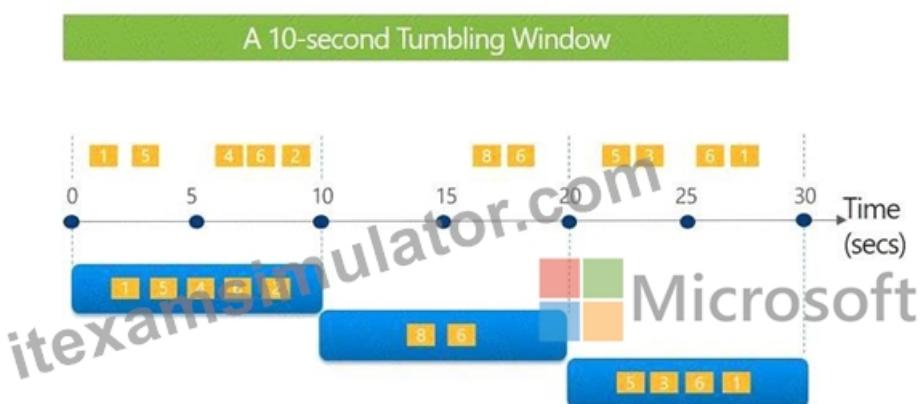
TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.

Analytic Function Syntax:

TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT DURATION clause>

[<WHEN clause>]) Box 2: Tumbling(minute 5) Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

**Tell me the count of Tweets per time zone every 10 seconds**



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second, 10)
```

Reference:

<https://docs.microsoft.com/en-us/stream-analytics-query/topone-azure-stream-analytics>

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun>

## NEW QUESTION # 192

• • • • •

**DP-300 Valid Exam Camp:** <https://www.itexamsimulator.com/DP-300-brain-dumps.html>

P.S. Free & New DP-300 dumps are available on Google Drive shared by ITEXamSimulator: [https://drive.google.com/open?id=1T96-wqcWGqYjhE-V6\\_1zdS45qp8x35P](https://drive.google.com/open?id=1T96-wqcWGqYjhE-V6_1zdS45qp8x35P)