



DP-300^{Q&As}

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/dp-300.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



**QUESTION 1**

You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently.

You enable columnstore archival compression.

What are two possible results of the configuration? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Queries that use the index will consume more disk I/O.
- B. Queries that use the index will retrieve fewer data pages.
- C. The index will consume more disk space.
- D. The index will consume more memory.
- E. Queries that use the index will consume more CPU resources.

Correct Answer: BE

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. In addition to saving space, data compression can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk.

Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data.

QUESTION 2

You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database.

Which database option should set?

- A. PARAMETERIZATION to FORCED
- B. PARAMETERIZATION to SIMPLE
- C. Delayed Durability to Forced
- D. READ_COMMITTED_SNAPSHOT to ON

Correct Answer: D

In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:

1.



The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.

2.

The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Incorrect Answers:

A: When the PARAMETERIZATION database option is set to SIMPLE, the SQL Server query optimizer may choose to parameterize the queries. This means that any literal values that are contained in a query are substituted with parameters. This process is referred to as simple parameterization. When SIMPLE parameterization is in effect, you cannot control which queries are parameterized and which queries are not.

B: You can specify that all queries in a database be parameterized by setting the PARAMETERIZATION database option to FORCED. This process is referred to as forced parameterization.

C: Delayed transaction durability is accomplished using asynchronous log writes to disk. Transaction log records are kept in a buffer and written to disk when the buffer fills or a buffer flushing event takes place. Delayed transaction durability reduces both latency and contention within the system.

Some of the cases in which you could benefit from using delayed transaction durability are:

1.

You can tolerate some data loss.

2.

You are experiencing a bottleneck on transaction log writes.

3.

Your workloads have a high contention rate.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql>

QUESTION 3

DRAG DROP

You create an Azure SQL managed instance and a job that performs backups.

You need to configure the job to notify a distribution group by email when the job fails. The solution must minimize administrative effort.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:



Actions Commands Cmdlets Statements

- Configure Database Mail.
- Configure a job notification.
- Configure an alert.
- Configure SendGrid.
- Create an operator.

Answer Area

Correct Answer:

Actions Commands Cmdlets Statements

-
-
- Configure an alert.
- Configure SendGrid.
-

Answer Area

Automate management tasks using SQL Agent jobs in Azure SQL Managed Instance

Using SQL Server Agent in SQL Server and SQL Managed Instance, you can create and schedule jobs that could be periodically executed against one or many databases to run Transact-SQL (T-SQL) queries and perform maintenance tasks.

Step 1: Configure Database Mail

If it isn't already enabled, first you would need to configure the Database Mail feature on SQL Managed Instance:

GO

```
EXEC sp_configure 'show advanced options', 1;
```

GO



```
RECONFIGURE;
```

```
GO
```

```
EXEC sp_configure '\\Database Mail XPs\\', 1;
```

```
GO
```

```
RECONFIGURE
```

Step 2: Create an operator

You can notify the operator that something happened with your SQL Agent jobs. An operator defines contact information for an individual responsible for the maintenance of one or more instances in SQL Managed Instance. Sometimes, operator responsibilities are assigned to one individual.

You can create operators using SQL Server Management Studio (SSMS) or the Transact-SQL script shown in the following example:

```
EXEC msdb.dbo.sp_add_operator @name=N'AzureSQLTeam', @enabled=1,  
@email_address=N'AzureSQLTeamn@contoso.com';
```

Step 3: Configure a job notification

Job notifications

SQL Agent jobs enable you to get notifications when the job finishes successfully or fails. You can receive notifications via email.

Example:

You can then modify any SQL Agent job and assign operators that will be notified via email if the job completes, fails, or succeeds using SSMS or the following T-SQL script:

```
EXEC msdb.dbo.sp_update_job @job_name=N'Load data using SSIS', @notify_level_email=3, -- Options are: 1 on  
succeed, 2 on failure, 3 on complete @notify_email_operator_name=N'AzureSQLTeam';
```

Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance>

QUESTION 4

HOTSPOT

You have an Azure subscription that contains an Azure SQL database.

The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource_semaphore waits.

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

```
SELECT
```

	▼
is_user_process	
wait_time	
wait_type	

```
    SUM(wait_time) AS total_wait_time_ms
```

```
FROM sys.
```

	▼
dm_exec_query_stats	
dm_exec_requests	
query_store_query	

```
    JOIN sys.dm_exec_sessions AS dmvs2
```

```
        ON dmvs1.session_id = dmvs2.session_id
```

```
WHERE is_user_process = 1
```

```
GROUP BY wait_type
```

```
ORDER BY SUM(wait_time) DESC;
```

Correct Answer:



Answer Area

```
SELECT
```

	▼
is_user_process	
wait_time	
wait_type	

```
    SUM(wait_time) AS total_wait_time_ms
```

```
FROM sys.
```

	▼
dm_exec_query_stats	
dm_exec_requests	
query_store_query	

```
    JOIN sys.dm_exec_sessions AS dmvs2
```

```
        ON dmvs1.session_id = dmvs2.session_id
```

```
WHERE is_user_process = 1
```

```
GROUP BY wait_type
```

```
ORDER BY SUM(wait_time) DESC;
```

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

QUESTION 5

HOTSPOT

You configure backup for an Azure SQL database as shown in the following exhibit.



Point-in-time-restore

Specify how long you want to keep your point-in-time backups. [Learn more](#)

How many days would you like PITR backups to be kept? ⓘ

14

Long-term retention

Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more](#)

Weekly LTR Backups

Keep weekly backups for:

Monthly LTR Backups

Keep the first backup of each month for:

Yearly LTR Backups

Keep an annual backup for:

Which weekly backup of the year would you like to keep?

Use the drop-down menus to select the answer choice the completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

Hot Area:



To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

▼
point-time restore (PITR) backup
yearly long-term retention (LTR) backup
weekly long-term retention (LTR) backup
monthly long-term retention (LTR) backup

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

▼
1 backup copy
52 backup copies
64 backup copies
65 backup copies

Correct Answer:



To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

▼
point-time restore (PITR) backup
yearly long-term retention (LTR) backup
weekly long-term retention (LTR) backup
monthly long-term retention (LTR) backup

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

▼
1 backup copy
52 backup copies
64 backup copies
65 backup copies