

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.leads4pass.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 managed instance.

You need to configure the SQL Server Agent service to email job notifications.

Which statement should you execute?

- A. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sysadmin_dbmail_profile';
- B. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'application_dbmail_profile';
- C. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'AzureManagedInstance_dbmail_profile';
- D. EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sys_dbmail_profile';

Correct Answer: C

-- Create a Database Mail profile

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = '\\AzureManagedInstance_dbmail_profile\\',  
@description = '\\E-mail profile used for messages sent by Managed Instance SQL Agent.\\';
```

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

QUESTION 2

You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.

Data to be loaded is identified by a column named LastUpdatedDate in the source table.

You plan to execute the pipeline every four hours.

You need to ensure that the pipeline execution meets the following requirements:

1.

Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits.

2.

Supports backfilling existing data in the table. Which type of trigger should you use?

- A. tumbling window
- B. on-demand
- C. event
- D. schedule

Correct Answer: A

The Tumbling window trigger supports backfill scenarios. Pipeline runs can be scheduled for windows in the past.
Incorrect Answers:

D: Schedule trigger does not support backfill scenarios. Pipeline runs can be executed only on time periods from the current time and the future.

Reference: <https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers>

QUESTION 3

DRAG DROP

You have SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

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.

Select and Place:

| Actions | Answer Area |
|--|-------------|
| Export a built-in policy. | |
| Create a custom policy based on a condition. | |
| Create a custom condition based on a built-in facet. | ⏪ |
| View the policy history. | ⏩ |
| Import a policy file. | |
| Run a policy evaluation. | |

Correct Answer:

Actions**Answer Area** Run a policy evaluation.

Reference: <https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-based-management/>

QUESTION 4

You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:

1.

Support failovers that do not require client applications to change their connection strings.

2.

Replicate the database to a secondary Azure region.

3.

Support failover to the secondary region. What should you include in the recommendation?

- A. failover groups
- B. transactional replication
- C. Availability Zones
- D. geo-replication

Correct Answer: A

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview>

QUESTION 5

You have SQL Server on an Azure virtual machine that contains a database named DB1.

You have an application that queries DB1 to generate a sales report.

You need to see the parameter values from the last time the query was executed.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable Last_Query_Plan_Stats in the master database
- B. Enable Lightweight_Query_Profiling in DB1
- C. Enable Last_Query_Plan_Stats in DB1
- D. Enable Lightweight_Query_Profiling in the master database
- E. Enable PARAMETER_SNIFFING in DB1

Correct Answer: BC

Last_Query_Plan_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm_exec_query_plan_stats.

Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT_QUERY_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;.

Incorrect Answers:

A: Enable it for DB1, not for the master database.

E: Parameter sensitivity, also known as "parameter sniffing", refers to a process whereby SQL Server "sniffs" the current parameter values during compilation or recompilation, and passes it along to the Query Optimizer so that they can be used to generate potentially more efficient query execution plans.

Parameter values are sniffed during compilation or recompilation for the following types of batches: Stored procedures
Queries submitted via sp_executesql Prepared queries

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>
<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>

[Latest DP-300 Dumps](#)

[DP-300 Practice Test](#)

[DP-300 Study Guide](#)