



# DP-300<sup>Q&As</sup>

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.passapply.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

Which counter should you monitor for real-time processing to meet the technical requirements?

- A. SU% Utilization
- B. CPU% utilization
- C. Concurrent users
- D. Data Conversion Errors

Correct Answer: B

Scenario: Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

To monitor the performance of a database in Azure SQL Database and Azure SQL Managed Instance, start by monitoring the CPU and IO resources used by your workload relative to the level of database performance you chose in selecting a particular service tier and performance level.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/monitor-tune-overview>

---

### QUESTION 2

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL. Which switch should you use to switch between languages?

- A. \[\]
- B. %
- C. \[\]
- D. @

Correct Answer: B

You can override the default language by specifying the language magic command % at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

Reference: <https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use>

---

### QUESTION 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure SQL database named Sales.



You need to implement disaster recovery for Sales to meet the following requirements:

1.

During normal operations, provide at least two readable copies of Sales.

2.

Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

---

#### QUESTION 4

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state.

What should you do?

A. Double the value of Data Flush interval

B. Decrease by half the value of Data Flush Interval

C. Double the value of Statistics Collection Interval

D. Decrease by half the value of Statistics Collection interval

Correct Answer: B

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the maximum size limit between storage size checks, it transitions to read-only mode.

Incorrect Answers:



C: Statistics Collection Interval: Defines the level of granularity for the collected runtime statistic, expressed in minutes. The default is 60 minutes. Consider using a lower value if you require finer granularity or less time to detect and mitigate issues. Keep in mind that the value directly affects the size of Query Store data.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

---

## QUESTION 5

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.

You need to resolve the performance issues.

What should you do?

- A. Implement memory-optimized tables.
- B. Run the DBCC FLUSHPROCINDBcommand.
- C. Replace the sequential index keys with nonsequential keys.
- D. Run the DBCC DBREINDEXcommand.

Correct Answer: A

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb-contention>

[Latest DP-300 Dumps](#)

[DP-300 PDF Dumps](#)

[DP-300 Braindumps](#)