



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.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

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

1.

Send the output to an Azure Synapse.

2.

Identify spikes and dips in time series data.

3.

Minimize development and configuration effort. Which should you include in the solution?

A. Azure SQL Database

B. Azure Databricks

C. Azure Stream Analytics

Correct Answer: C

Anomalies can be identified by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics

Reference: <https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/>
<https://docs.microsoft.com/en-us/azure/stream-analytics/azure-synapse-analytics-output>

QUESTION 2

Which audit log destination should you use to meet the monitoring requirements?

A. Azure Storage

B. Azure Event Hubs

C. Azure Log Analytics

Correct Answer: C

Scenario: Use a single dashboard to review security and audit data for all the PaaS databases.

With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics.

Note: Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/visualize/tutorial-logs-dashboards>



QUESTION 3

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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query `sys.dm_exec_requests` and discover that the wait type is `PAGELATCH_UP` and the `wait_resource` is `2:3:905856`.

You need to improve system performance.

Solution: You shrink the transaction log file.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

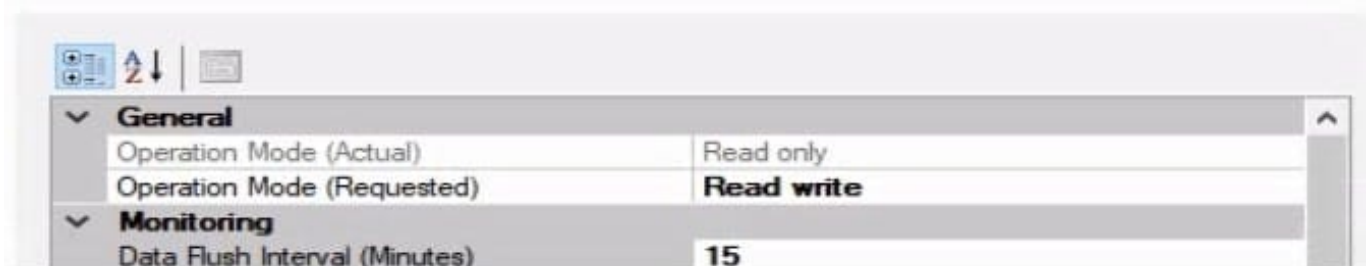
Reference: <https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

QUESTION 4

HOTSPOT

You have a database on a SQL Server on Azure Virtual Machines instance.

The current state of Query Store for the database is shown in the following exhibit.



Hot Area:



Query Store will retain [answer choice] queries for evaluation.

a selective set of
Place Holder
Place Holder
Place Holder

To change Operation Mode(Actual) to Read write without losing any data, you must modify the [answer choice] setting.

Max Size(MB)
Place Holder
Place Holder
Place Holder


Correct Answer:



Query Store will retain [answer choice] queries for evaluation.



To change Operation Mode(Actual) to Read write without losing any data, you must modify the [answer choice] setting.



QUESTION 5

A company plans to use Apache Spark analytics to analyze intrusion detection data.

You need to recommend a solution to analyze network and system activity data for malicious activities and policy violations. The solution must minimize administrative efforts.

What should you recommend?

- A. Azure Data Lake Storage
- B. Azure Databricks
- C. Azure HDInsight
- D. Azure Data Factory

Correct Answer: C

Azure HDInsight offers pre-made, monitoring dashboards in the form of solutions that can be used to monitor the workloads running on your clusters. There are solutions for Apache Spark, Hadoop, Apache Kafka, live long and process (LLAP), Apache HBase, and Apache Storm available in the Azure Marketplace.



Note: With Azure HDInsight you can set up Azure Monitor alerts that will trigger when the value of a metric or the results of a query meet certain conditions. You can condition on a query returning a record with a value that is greater than or less than a certain threshold, or even on the number of results returned by a query. For example, you could create an alert to send an email if a Spark job fails or if a Kafka disk usage becomes over 90 percent full.

Reference: <https://azure.microsoft.com/en-us/blog/monitoring-on-azure-hdinsight-part-4-workload-metrics-and-logs/>

QUESTION 6

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4- TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A. distributed availability groups
- B. database mirroring
- C. log shipping
- D. Database Migration Assistant

Correct Answer: D

Ref: <https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-azure-sql>

QUESTION 7

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 8

HOTSPOT

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

```
SELECT
    payment_type,
    SUM(fare_amount) AS fare_total
FROM OPENROWSET (
    BULK 'csv/busfare/tripdata_2020*.csv',
    DATA_SOURCE = 'BusData',
    FORMAT = 'CSV', PARSER_VERSION = '2.0',
    FIRSTROW - 2
)
WITH (
    payment_type INT 10,
    fare_amount FLOAT 11
) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

Hot Area:



Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

▼
CSV files in the tripdata_2020 subfolder
files that have files names beginning with "tripdata_2020"
CSV files that have file names containing "tripdata_202"
CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.

▼
a header
a data
an empty

Correct Answer:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

▼
CSV files in the tripdata_2020 subfolder
files that have files names beginning with "tripdata_2020"
CSV files that have file names containing "tripdata_202"
CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.

▼
a header
a data
an empty

Box 1: CSV files that have file named beginning with "tripdata_2020"

Box 2: a header

FIRSTROW = \\first_row\\

Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset>

QUESTION 9

DRAG DROP

You have an Azure subscription that contains the resources shown in the following table.



Name	Type	Description
SQL1	SQL Server on Azure Virtual Machines	Not applicable
db1	Microsoft SQL Server database	Hosted on SQL1
mysqlbackups	General purpose v2 storage account	Not applicable

You need to back up db1 to mysqlbackups, and then restore the backup to a new database named db2 that is hosted on SQL1. The solution must ensure that db1 is backed up to a stripe set.

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

Select and Place:



```
RESTORE DATABASE db2 FROM URL = URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_4.bak'  
WITH CREDENTIAL = 'sqlbackup', RECOVERY,  
MOVE 'db1_mdf' TO  
'D:\Data\db2_mdf.mdf',  
MOVE 'db1_log' TO  
'D:\Logs\db2_log.ldf'
```

1

```
RESTORE DATABASE db2 FROM URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_4.bak'  
WITH RECOVERY,  
MOVE 'db1_mdf' TO  
'D:\Data\db2_mdf.mdf',  
MOVE 'db1_log' TO  
'D:\Logs\db2_log.ldf'
```

2

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY  
=  
'sqlsamplebackup'  
, SECRET = '<mystorageaccountaccesskey>';  
GO
```

3

```
BACKUP DATABASE db1  
TO URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_4.bak'  
WITH CREDENTIAL = 'sqlbackup';  
GO
```

4

```
BACKUP DATABASE db1  
TO URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net/  
/backups/db1_4.bak'  
GO
```

5

```
CREATE CREDENTIAL  
[https://mysqlbackups.blob.core.windows.net/  
/backups]  
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',  
SECRET = '<SAS_TOKEN>'  
GO
```

6

Answer Area

A B C

Correct Answer:



```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup', RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

1

```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
WITH RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

2

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY =
'sqlsamplebackup'
, SECRET = '<mystorageaccountaccesskey>;'
GO
```

3

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup';
GO
```

4

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net/backups/db1_4.bak'
GO
```

5

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

6

Answer Area

6

5

2

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-server-ver15>

QUESTION 10

You have 20 Azure SQL databases provisioned by using the vCore purchasing model.

You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct



answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. total size of all the databases
- B. geo-replication support
- C. number of concurrently peaking databases * peak CPU utilization per database
- D. maximum number of concurrent sessions for all the databases
- E. total number of databases * average CPU utilization per database

Correct Answer: ACE

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model: MAX(,