

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.pass2lead.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 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 2

You have an Azure SQL database named DB1.

You have a table name Table1 that has 20 columns of type CHAR(400). Row compression for Table1 is enabled.

During a database audit, you discover that none of the fields contain more than 150 characters.

You need to ensure that you can apply page compression to Table1.

What should you do?

A. Configure the columns as sparse.

B. Change the column type to NVARCHAR(MAX).

C. Change the column type to VARCHAR(MAX).

D. Change the column type to VARCHAR(200).

Correct Answer: D

Reference:

<https://www.sqlshack.com/sql-varchar-data-type-deep-dive/>

<https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/>

QUESTION 3

You have an Azure Stream Analytics job.

You need to ensure that the job has enough streaming units provisioned.

You configure monitoring of the SU % Utilization metric.

Which two additional metrics should you monitor? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Late Input Events

B. Out of order Events

C. Backlogged Input Events

D. Watermark Delay

E. Function Events

Correct Answer: CD

To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact.

Note: Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently nonzero, you should scale out your job, by increasing the SUs.

Reference: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring>

QUESTION 4

You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1.

You need to configure SQL1 to use mixed mode authentication.

Which procedure should you run?

- A. sp_addremotelogin
- B. xp_instance_regwrite
- C. sp_change_users_login
- D. xp_grant_login

Correct Answer: B

Change authentication mode (Transact-SQL)

The following example changes Server Authentication from mixed mode (Windows and SQL) to Windows only.

```
USE [master]
```

```
GO
```

```
EXEC xp_instance_regwrite N\\HKEY_LOCAL_MACHINE\\,
```

```
N\\Software\\Microsoft\\MSSQLServer\\MSSQLServer\\, N\\'LoginMode\\', REG_DWORD, 1; GO
```

Reference: <https://learn.microsoft.com/en-us/sql/database-engine/configure-windows/change-server-authentication-mode>

QUESTION 5

HOTSPOT

You are building a database in an Azure Synapse Analytics serverless SQL pool.

You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container.

Records are structured as shown in the following sample.

```
{  
  "id":123,  
  "address_housenumber": "19c",  
  "address_line1": "Memory Lane",  
  "applicant1_name": "Jane",  
  "applicant2_name": "Dev"  
}
```

The records contain two applicants at most.

You need to build a table that includes only the address fields.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
▼ applications
CREATE EXTERNAL TABLE
CREATE TABLE
CREATE VIEW
WITH (
    LOCATION = 'applications/',
    DATA_SOURCE = applications_ds,
    FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
▼ (BULK 'https://contoso1.dfs.core.windows.net/
applications/year=*/*.parquet',
CROSS APPLY
OPENJSON
OPENROWSET
    FORMAT = 'PARQUET') AS [r]
GO
```

Correct Answer:

Answer Area

```
▼ applications
CREATE EXTERNAL TABLE
CREATE TABLE
CREATE VIEW
WITH (
    LOCATION = 'applications/',
    DATA_SOURCE = applications_ds,
    FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
▼ (BULK 'https://contoso1.dfs.core.windows.net/
applications/year=*/*.parquet',
CROSS APPLY
OPENJSON
OPENROWSET
    FORMAT = 'PARQUET') AS [r]
GO
```

Box 1: CREATE EXTERNAL TABLE An external table points to data located in Hadoop, Azure Storage blob, or Azure Data Lake Storage. External tables are used to read data from files or write data to files in Azure Storage. With Synapse SQL, you can use external tables to read external data using dedicated SQL pool or serverless SQL pool.

Syntax:

```
CREATE EXTERNAL TABLE { database_name.schema_name.table_name | schema_name.table_name | table_name }
([,...n]) WITH (
    LOCATION = '\\folder_or_filepath\\',
    DATA_SOURCE = external_data_source_name,
    FILE_FORMAT = external_file_format_name
```

Box 2. OPENROWSET

When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.

Example:

```
AS
SELECT decennialTime, stateName, SUM(population) AS population
FROM
```

```
OPENROWSET(BULK '\\https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*/*.parquet\\',
```

```
FORMAT=\\'PARQUET\\') AS [r] GROUP BY decennialTime, stateName GO
```

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

QUESTION 6

You have an Azure SQL database.

You discover that the plan cache is full of compiled plans that were used only once.

You run the select * from sys.database_scoped_configurations Transact-SQL command and receive the results shown in the following table.

configuration_id	name	value	is_value_default
1	LEGACY_CARDINALITY_ESTIMATION	0	1
2	QUERY_OPTIMIZER_HOTFIXES	0	1
3	OPTIMIZE_FOR_AD_HOC_WORKLOADS	0	1
4	ACCELERATED_PLAN_FORCING	1	1

You need relieve the memory pressure. What should you configure?

- A. LEGACY_CARDINALITY_ESTIMATION
- B. QUERY_OPTIMIZER_HOTFIXES
- C. OPTIMIZE_FOR_AD_HOC_WORKLOADS
- D. ACCELERATED_PLAN_FORCING

Correct Answer: C

OPTIMIZE_FOR_AD_HOC_WORKLOADS = { ON | OFF }

Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE_FOR_AD_HOC_WORKLOADS is enabled for a database, a

compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.

Incorrect Answers:

A: LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY } Enables you to set the query optimizer cardinality estimation model to the SQL Server 2012 and earlier version independent of the compatibility level of the database. The default is OFF, which sets the query optimizer cardinality estimation model based on the compatibility level of the database.

B: QUERY_OPTIMIZER_HOTFIXES = { ON | OFF | PRIMARY }

Enables or disables query optimization hotfixes regardless of the compatibility level of the database. The default is OFF,

which disables query optimization hotfixes that were released after the highest available compatibility level was introduced for a specific version (post-RTM).

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql>

QUESTION 7

You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information.

You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information.

What should you include in the recommendation?

- A. row-level security
- B. data masking
- C. Always Encrypted
- D. column-level security

Correct Answer: B

Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card.

Example:

XXXX-XXXX-XXXX-1234

QUESTION 8

HOTSPOT

You have an Azure SQL database named DB 1 in the General Purpose service tier.

You need to monitor DB 1 by using SQL Insights.

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

NOTE: Each correct selection is worth one point.

Hot Area:

To collect monitoring data, use: ▼

- A virtual machine
- An Azure function
- The Azure Monitor agent

To store monitoring data, create: ▼

- A Log Analytics workspace
- An Azure SQL database
- An Azure Storage account

Correct Answer:

To collect monitoring data, use: ▼

- A virtual machine
- An Azure function
- The Azure Monitor agent**

To store monitoring data, create: ▼

- A Log Analytics workspace
- An Azure SQL database**
- An Azure Storage account

Box 1 = Azure Monitor Agent Box 2 = An Azure SQL database <https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql>

QUESTION 9

HOTSPOT

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

Name	Type	Configuration
DB1	Azure SQL Database	Hyperscale service tier No secondary replicas
App1	Azure Web Apps	App1 has read-only access to DB1. There are multiple instances of App1.

You need to create a read-only replica of DB1 and configure the App1 instances to use the replica.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To add read-only replicas of DB1:

▼
Create a replica on the same logical server.
Create a new logical server and configure geo-replication.
Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

▼
Add an ApplicationIntent entry to the connection string.
Add a MultiSubnetFailover entry to the App1 connection string.
Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

Correct Answer:

Answer Area

To add read-only replicas of DB1:

▼
Create a replica on the same logical server.
Create a new logical server and configure geo-replication.
Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

▼
Add an ApplicationIntent entry to the connection string.
Add a MultiSubnetFailover entry to the App1 connection string.
Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

Reference: <https://sqlserverguides.com/read-only-replica-azure-sql/>

QUESTION 10

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

CREATE TABLE dbo.Employee

```
(  
    EmployeeID          INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,  
    FirstName           VARCHAR(100) NOT NULL,  
    LastName            VARCHAR(100) NOT NULL,  
    Title               VARCHAR(100) NULL,  
    LastHireDate        DATETIME NULL,  
    StreetAddress1      VARCHAR(500) NOT NULL,  
    StreetAddress2      VARCHAR(500) NOT NULL,  
    StreetAddress3      VARCHAR(500) NOT NULL,  
    City                VARCHAR(200) NOT NULL,  
    StateName           VARCHAR(20) NOT NULL,  
    Salary              VARCHAR(20) NULL,  
    PhoneNumber         VARCHAR(20) NOT NULL  
)
```

You need to recommend changes to the data types to reduce storage and improve performance. Which two actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Change Salary to the money data type.
- B. Change PhoneNumber to the float data type.
- C. Change LastHireDate to the datetime2(7) data type.
- D. Change PhoneNumber to the bigint data type.
- E. Change LastHireDate to the date data type.

Correct Answer: AE

QUESTION 11

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

Name	Type	Azure region
VM1	Azure virtual machine	West US 2
MI1	Azure SQL Managed Instance	East US

You need to configure a connection between VM1 and MI1. The solution must meet the following requirements:

1.
The connection must be encrypted.
2.
Network latency must be minimized. What should you implement?

- A. a site-to-site VPN
- B. virtual network peering
- C. private endpoints
- D. service endpoints

Correct Answer: C

QUESTION 12

You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases.

The databases perform slower than expected.

You need to identify whether the performance issue relates to the use of tempdb by Azure SQL databases on sqlsrv1.

What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

Correct Answer: D

Microsoft Azure SQL Database and Azure SQL Managed Instance enable a subset of dynamic management views to diagnose performance problems, which might be caused by blocked or long-running queries, resource bottlenecks, poor

query plans, and so on.

This include edentifying tempdb performance issues:

When identifying IO performance issues, the top wait types associated with tempdb issues is PAGELATCH_* (not PAGEIOLATCH_*).

Reference:

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

QUESTION 13

You are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.

You need to create a surrogate key for the table. The solution must provide the fastest query performance.

What should you use for the surrogate key?

- A. an IDENTITY column
- B. a GUID column
- C. a sequence object

Correct Answer: A

Dedicated SQL pool supports many, but not all, of the table features offered by other databases. Surrogate keys are not supported. Implement it with an Identity column. Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables- overview>

QUESTION 14

You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:

1.
Support real-time data replication to a different geographic region.

2.
Use Azure as a disaster recovery target.

3.
Minimize costs and administrative effort.

What should you include in the recommendation?

- A. database mirroring on an instance of SQL Server on Azure Virtual Machines
- B. availability groups for SQL Server on Azure Virtual Machines
- C. an Azure SQL Managed Instance link
- D. transactional replication to an Azure SQL Managed Instance

Correct Answer: D

Transactional replication is a feature of Azure SQL Managed Instance and SQL Server that enables you to replicate data from a table in Azure SQL Managed Instance or a SQL Server instance to tables placed on remote databases. This feature allows you to synchronize multiple tables in different databases.D

You can use transactional replication to push changes made in an Azure SQL Managed Instance to:

- 1.
A SQL Server database - on-premises or on Azure VM
- 2.
A database in Azure SQL Database
- 3.
An instance database in Azure SQL Managed Instance

Incorrect:

Not A: Operates within the same location.

Database mirroring is a solution for increasing the availability of a SQL Server database.

Database mirroring maintains two copies of a single database that must reside on different server instances of SQL Server Database Engine. Typically, these server instances reside on computers in different locations.

Not B: Always On availability group on SQL Server on Azure VMs - only within a single region

Always On availability groups on Azure Virtual Machines are similar to Always On availability groups on-premises, and rely on the underlying Windows Server Failover Cluster.

VM redundancy

To increase redundancy and high availability, SQL Server VMs should either be in the same availability set, or different availability zones.

Placing a set of VMs in the same availability set protects from outages within a data center caused by equipment failure (VMs within an Availability Set do not share resources) or from updates (VMs within an availability set are not updated at

the same time).

Availability Zones protect against the failure of an entire data center, with each Zone representing a set of data centers within a region.

Reference: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/replication-transactional-overview>
<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-overview>

QUESTION 15

You have an Azure subscription that contains an Azure SQL database named SQL1.

SQL1 is in an Azure region that does not support availability zones.

You need to ensure that you have a secondary replica of SQL1 in the same region.

What should you use?

- A. log shipping
- B. auto-failover groups
- C. active geo-replication
- D. Microsoft SQL Server failover clusters

Correct Answer: C

Active geo-replication is a feature that lets you to create a continuously synchronized readable secondary database for a primary database. The readable secondary database may be in the same Azure region as the primary, or, more commonly, in a different region. This kind of readable secondary databases are also known as geo-secondaries, or geo-replicas.

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

[DP-300 PDF Dumps](#)

[DP-300 Study Guide](#)

[DP-300 Exam Questions](#)