

## Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

<https://www.2passeasy.com/dumps/DP-300/>



### NEW QUESTION 1

- (Exam Topic 5)

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales. You need to initiate a backup of Sales.

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.

BACKUP DATABASE Sales

	▼
TO DISK = \\BackupSystem\BackupDisk1\Sales.bak'	
TO DISK = 'X:\BAK\Sales.bak'	
TO 'Sales_Backup'	
TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'	

WITH STATS = 5,

	▼
WITH COPY_ONLY;	
WITH ENCRYPTION;	
WITH FILE_SNAPSHOT;	
WITH NO_TRUNCATE	

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak' Native database backup in Azure SQL Managed Instance.

You can backup any database using standard BACKUP T-SQL command: BACKUP DATABASE tpcc2501

TO URL = 'https://myacc.blob.core.windows.net/testcontainer/tpcc2501.bak'

WITH COPY\_ONLY

Box 2: WITH COPY\_ONLY

Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-insta>

### NEW QUESTION 2

- (Exam Topic 5)

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 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

**Answer:** D

#### Explanation:

The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting.

You can use the Intelligent Insights performance diagnostics log of Azure SQL Database to troubleshoot performance issues.

Reference:

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

### NEW QUESTION 3

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.

You need to protect the contents of Column1 at rest, in transit, and in use.

How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Encryption key:

Encryption type:

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Column encryption Key

Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

**NEW QUESTION 4**

- (Exam Topic 5)

You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables. Which distribution type should you recommend to minimize data movement?

- A. HASH
- B. REPLICATE
- C. ROUND\_ROBIN

**Answer:** B

**Explanation:**

A replicated table has a full copy of the table available on every Compute node. Queries run fast on replicated tables since joins on replicated tables don't require data movement. Replication requires extra storage, though, and isn't practical for large tables.

Reference:

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

**NEW QUESTION 5**

- (Exam Topic 5)

You have an Azure Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

You need to recommend a solution for account1 that meets the following requirements:

- > Automatically deletes the logs at the end of each retention period
- > Minimizes storage costs

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

NOTE: Each correct selection is worth one point.

To minimize storage costs:

To delete the logs automatically:

- A. Mastered

B. Not Mastered

**Answer:** A

**Explanation:**

A picture containing text Description automatically generated

Box 1: Store the infrastructure logs in the Cool access tier the application logs in the Archive access tier Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Azure Blob storage lifecycle management rules

Blob storage lifecycle management offers a rich, rule-based policy that you can use to transition your data to the best access tier and to expire data at the end of its lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**NEW QUESTION 6**

- (Exam Topic 5)

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

**Answer:** C

**Explanation:**

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.

Reference:

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

**NEW QUESTION 7**

- (Exam Topic 5)

You have a 50-TB Microsoft SQL Server database named DB1.

You need to reduce the time it takes to perform database consistency checks of DB1.

Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

DBCC CHECKDB ([DB1], NOINDEX  
REPAIR\_FAST  
REPAIR\_REBUILD with ALL\_ERRORMSG  
NO\_INFOMSGS  
PHYSICAL\_ONLY

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Table Description automatically generated with low confidence

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-ser>

**NEW QUESTION 8**

- (Exam Topic 5)

You configure version control for an Azure Data Factory instance as shown in the following exhibit.



Connections

Linked services

Integration runtimes

Azure Purview (Preview)

Source control

Git configuration

ARM template

Parameterization template

Author

Triggers

Global parameters

Security

Customer managed key

Managed private endpoints

Git repository

Git repository information associated with your data factory. [CI/CD best practices](#)

Setting

Disconnect

Repository type

Azure DevOps Git

Azure DevOps Account

CONTOSO

Project name

Data

Repository name

dwh\_batchetl

Collaboration branch

main

Publish branch

adf\_publish

Root folder

/

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

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

/

adf\_publish

main

Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

/contososales

/dwh\_batchetl/adf\_publish/contososales

/main

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: adf\_publish

By default, data factory generates the Resource Manager templates of the published factory and saves them into a branch called adf\_publish. To configure a custom publish branch, add a publish\_config.json file to the root folder in the collaboration branch. When publishing, ADF reads this file, looks for the field publishBranch, and saves all Resource Manager templates to the specified location. If the branch doesn't exist, data factory will automatically create it. And example of what this file looks like is below:

```
{
  "publishBranch": "factory/adf_publish"
}
```

Box 2: /dwh\_barchlet/ adf\_publish/contososales

RepositoryName: Your Azure Repos code repository name. Azure Repos projects contain Git repositories to manage your source code as your project grows. You can create a new repository or use an existing repository that's already in your project.

Reference:

https://docs.microsoft.com/en-us/azure/data-factory/source-control

NEW QUESTION 9

- (Exam Topic 5)

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails. You need to configure Azure Synapse Analytics to receive the data.

Which four 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.

## Actions

## Answer Area

Create an external data source for Azure Blob storage.

Create a master key on database.

Enable Transparent Data Encryption.

Create the external table FactSalesOrderDetails.

Load the data to a staging table.

Create an external file format to map the parquet files.



- A. Mastered
- B. Not Mastered

**Answer:** A

### Explanation:

Graphical user interface, text, application, chat or text message Description automatically generated

To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table.

Step 1: Create a master key on database.

\* 1. Create a master key on the database. The master key is required to encrypt the credential secret. (Create a database scoped credential for Azure blob storage.)

Step 2: Create an external data source for Azure Blob storage.

\* 2. Create an external data source with CREATE EXTERNAL DATA SOURCE.. Step 3: Create an external file format to map the parquet files.

\* 3. Create an external file format with CREATE EXTERNAL FILE FORMAT. Step 4. Create an external table FactSalesOrderDetails

\* 4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE. Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage>

### NEW QUESTION 10

- (Exam Topic 5)

You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.

You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

- A. Upgrade workspace1 to the Premium pricing tier.
- B. Configure a global init script for workspace1.
- C. Create a pool in workspace1.
- D. Create a cluster policy in workspace1.

**Answer:** C

### Explanation:

You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.

Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.

Reference:

<https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html>

### NEW QUESTION 10

- (Exam Topic 5)

You have an Azure SQL managed instance.

You need to restore a database named DB1 by using Transact-SQL.

Which command should you run? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

RESTORE		DB1 FROM	
DATABASE			DISK = N'\\NAS01\\SQLBackups\\DB1.bak';
FILE			TAPE = N'\\Tape0'
LOG			URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated

**NEW QUESTION 14**

- (Exam Topic 5)

You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled. You need to create a new database user for app1.

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.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

<https://learn.microsoft.com/en-us/azure/app-service/tutorial-connect-msi-sql-database?tabs=windowsclient%2Ce>

**NEW QUESTION 15**

- (Exam Topic 5)

You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.

During peak usage, the database will require the following:

- > 24 cores
- > 500 GB of storage
- > 124 GB of memory
- > More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard. Which service tier should you use during peak usage?

- A. Business Critical
- B. Premium
- C. Hyperscale

**Answer:** A

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vcare-single-databases#business-critic>

**NEW QUESTION 17**

- (Exam Topic 5)

You have an Azure SQL database named DB1. You run a query while connected to DB1.

You review the actual execution plan for the query, and you add an index to a table referenced by the query. You need to compare the previous actual execution plan for the query to the Live Query Statistics.

What should you do first in Microsoft SQL Server Management Studio (SSMS)?

- A. For DB1, set QUERY\_CAPTURE\_MODE of Query Store to All.
- B. Run the SET SHOWPLAN\_ALL Transact-SQL statement.
- C. Save the actual execution plan.
- D. Enable Query Store for DB1.

**Answer:** C

**Explanation:**

The Plan Comparison menu option allows side-by-side comparison of two different execution plans, for easier identification of similarities and changes that explain the different behaviors for all the reasons stated above. This option can compare between:

Two previously saved execution plan files (.sqlplan extension).

One active execution plan and one previously saved query execution plan. Two selected query plans in Query Store.

**NEW QUESTION 21**

- (Exam Topic 5)

You have a new Azure SQL database. The database contains a column that stores confidential information. You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Turn on auditing and write audit logs to an Azure Storage account.



- B. Add extended properties to the column.
- C. Turn on Advanced Data Security for the Azure SQL server.
- D. Apply sensitivity labels named Highly Confidential to the column.
- E. Turn on Azure Advanced Threat Protection (ATP).

**Answer:** ACD

**Explanation:**

C: Advanced Data Security (ADS) is a unified package for advanced SQL security capabilities. ADS is available for Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics. It includes functionality for discovering and classifying sensitive data

D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes that have been added to the SQL Server database engine. This metadata can then be used for advanced, sensitivity-based auditing and protection scenarios.

A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called data\_sensitivity\_information. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	7.125	Microsoft SQL Server Management Studio - Query	41	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview

**NEW QUESTION 22**

- (Exam Topic 5)

You have the following Transact-SQL query.

```
SELECT
    [file_id] AS [File ID],
    [type] AS [File Type],
    substring([physical_name], 1,1) AS [Drive],
    [name] AS [Logical Name],
    [physical_name] AS [Physical Name],
    CAST([size] as DECIMAL(38,0))/128.0 AS [ColumnA],
    CAST(FILEPROPERTY([name], 'SpaceUsed') AS DECIMAL(38,0))/128.0 AS
[ColumnB],
    (CAST([size] AS DECIMAL(38,0))/128.0) - (CAST(FILEPROPERTY([name],
'SpaceUsed') AS DECIMAL (38,0))/128.0) AS [ColumnC],
    [max_size] AS [ColumnD],
    [is_percent_growth] AS [Percent Growth Enabled],
    [growth] AS [Growth Rate],
    SYSDATETIME() AS [Current Date]
FROM sys.database_files;
```

Which column returned by the query represents the free space in each file?

- A. ColumnA
- B. ColumnB
- C. ColumnC
- D. ColumnD

**Answer:** C

**Explanation:**

Example:

Free space for the file in the below query result set will be returned by the FreeSpaceMB column.

```
SELECT DB_NAME() AS DbName,
name AS FileName, type_desc,
size/128.0 AS CurrentSizeMB,
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB FROM sys.database_files
WHERE type IN (0,1);
```

Reference:

https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/

**NEW QUESTION 23**

- (Exam Topic 5)

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.

Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted. You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



## Values

## Answer Area

CustomerKey

HASH

ROUND\_ROBIN

REPLICATE

OrderDateKey

SalesOrderNumber

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = ( ) ([ProductKey])
, PARTITION ( [ ] RANGE RIGHT FOR VALUES
                (20170101, 20180101, 20190101, 20200101, 20210101)
            )
)
```

- A. Mastered
- B. Not Mastered

Answer: A

### Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: HASH

Box 2: OrderDateKey

In most cases, table partitions are created on a date column.

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management. For example, rather than execute a DELETE statement to delete all rows in a table where the order\_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse> <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool>

### NEW QUESTION 27

- (Exam Topic 5)

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.

You scale up the virtual machine to 16 vCPUSs and 64 GB of memory. You need to provide the lowest latency for tempdb.

What is the total number of data files that tempdb should contain?

- A. 2
- B. 4
- C. 8
- D. 64

Answer: D

### Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical

processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

### NEW QUESTION 31

- (Exam Topic 5)

You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in five minute intervals and report only events that arrive during the interval.

The output will be sent to a Delta Lake table. Which output mode should you use?

- A. complete
- B. append
- C. update

Answer: A

### Explanation:

Complete mode: You can use Structured Streaming to replace the entire table with every batch. Reference:

<https://docs.databricks.com/delta/delta-streaming.html>

NEW QUESTION 34

- (Exam Topic 5)

You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer\_Phone column. The mask must meet the following requirements:

- > The first six numerals of each customer’s phone number must be masked.
- > The last four digits of each customer’s phone number must be visible.
- > Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

Exposed Prefix:

▼

0

1

3

5

Padding String:

▼

X

XXXXXX

XXX-XXX

XXX-XXX-

x[3]-x[3]

Exposed Suffix:

▼

0

1

3

5

- A. Mastered  
B. Not Mastered

Answer: A

Explanation:

Box 1: 0

Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix

Box 2: xxx-xxx

Box 3: 5 Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>

NEW QUESTION 36

- (Exam Topic 5)

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.

▼ 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',
    FORMAT = 'PARQUET') AS [r]
GO
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

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 } ( <column\_definition> [ ,...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=\*/ FORMAT='PARQUET') AS [r] GROUP BY decennialTime, stateName GO

Reference:

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

NEW QUESTION 38

- (Exam Topic 5)

From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays. The data contains the following columns:

Name	Sample value
Date	15 Jan 2021
EventCategory	Videos
EventAction	Play
EventLabel	Contoso Promotional
ChannelGrouping	Social
TotalEvents	150
UniqueEvents	120
SessionsWithEvents	99

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension. To which table should you add each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



EventCategory:

ChannelGrouping:

TotalEvents:

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Graphical user interface, application, table Description automatically generated

Box 1: FactEvents

Fact tables store observations or events, and can be sales orders, stock balances, exchange rates, temperatures, etc.

Box 2: DimChannel

Dimension tables describe business entities – the things you model. Entities can include products, people, places, and concepts including time itself. The most consistent table you'll find in a star schema is a date dimension table. A dimension table contains a key column (or columns) that acts as a unique identifier, and descriptive columns.

Box 3: DimEvent Reference:

<https://docs.microsoft.com/en-us/power-bi/guidance/star-schema>

**NEW QUESTION 43**

- (Exam Topic 5)

You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1.

Several users execute adhoc queries to DW1 concurrently. You regularly perform automated data loads to DW1.

You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run.

What should you do?

- A. Assign a smaller resource class to the automated data load queries.
- B. Create sampled statistics to every column in each table of DW1.
- C. Assign a larger resource class to the automated data load queries.
- D. Hash distribute the large fact tables in DW1 before performing the automated data loads.

Answer: C

**Explanation:**

The performance capacity of a query is determined by the user's resource class.

Smaller resource classes reduce the maximum memory per query, but increase concurrency. Larger resource classes increase the maximum memory per query, but reduce concurrency. Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workloadman>

**NEW QUESTION 47**

- (Exam Topic 5)

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

- A. role assignments
- B. account keys
- C. shared access signatures (SAS)
- D. Azure Active Directory (Azure AD) identities

Answer: C

**Explanation:**

A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:

What resources the client may access.



What permissions they have to those resources. How long the SAS is valid.

Note: Data Lake Storage Gen2 supports the following authorization mechanisms:

- > Shared Key authorization
- > Shared access signature (SAS) authorization
- > Role-based access control (Azure RBAC)
- > Shared Key authorization
- > Shared access signature (SAS) authorization
- > Role-based access control (Azure RBAC)
- > Access control lists (ACL)

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

#### NEW QUESTION 49

- (Exam Topic 5)

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

**Answer: D**

#### Explanation:

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

- > The READ COMMITTED isolation level with the READ\_COMMITTED\_SNAPSHOT database option set to ON.
- > 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.

Reference:

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

#### NEW QUESTION 52

- (Exam Topic 5)

You have an Azure subscription.

You plan to deploy an Azure SQL database by using an Azure Resource Manager template.

How should you complete the template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

#### Answer Area

```
{
  "resources": [
    {
      "type": 
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
      ...
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",
          ...
          "dependsOn": [
            "properties": [
              "tags": [
                "[resourceId('Microsoft.Sql/servers', concat(parameters('name1')))]"
              ]
            ]
          ]
        }
      ]
    }
  ]
}
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated

Reference:

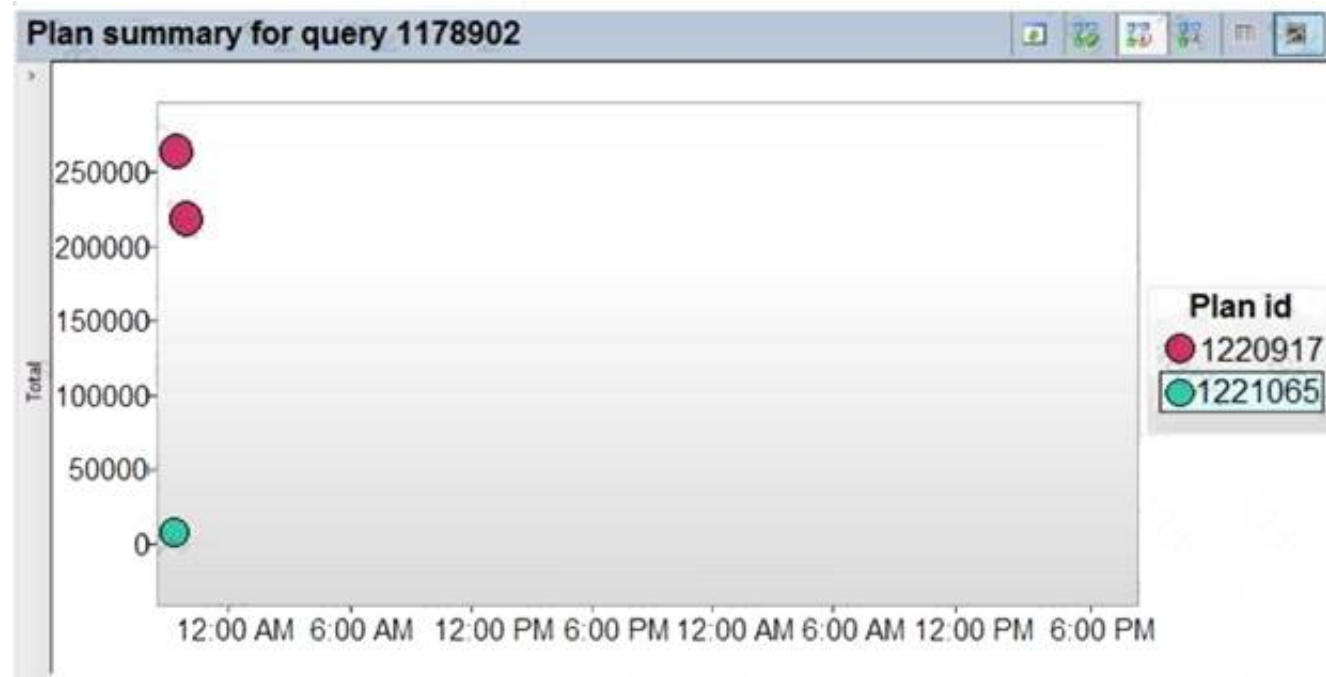
<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart>

**NEW QUESTION 54**

- (Exam Topic 5)

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

You view a plan summary that shows the duration in milliseconds of each execution of query 1178902 as shown in the following exhibit:



What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

- A. Force the query execution plan for plan 1221065.
- B. Run the DBCC FREEPROCCACHE command.
- C. Force the query execution plan for plan 1220917.
- D. Disable parameter sniffing.

**Answer:** C

**Explanation:**

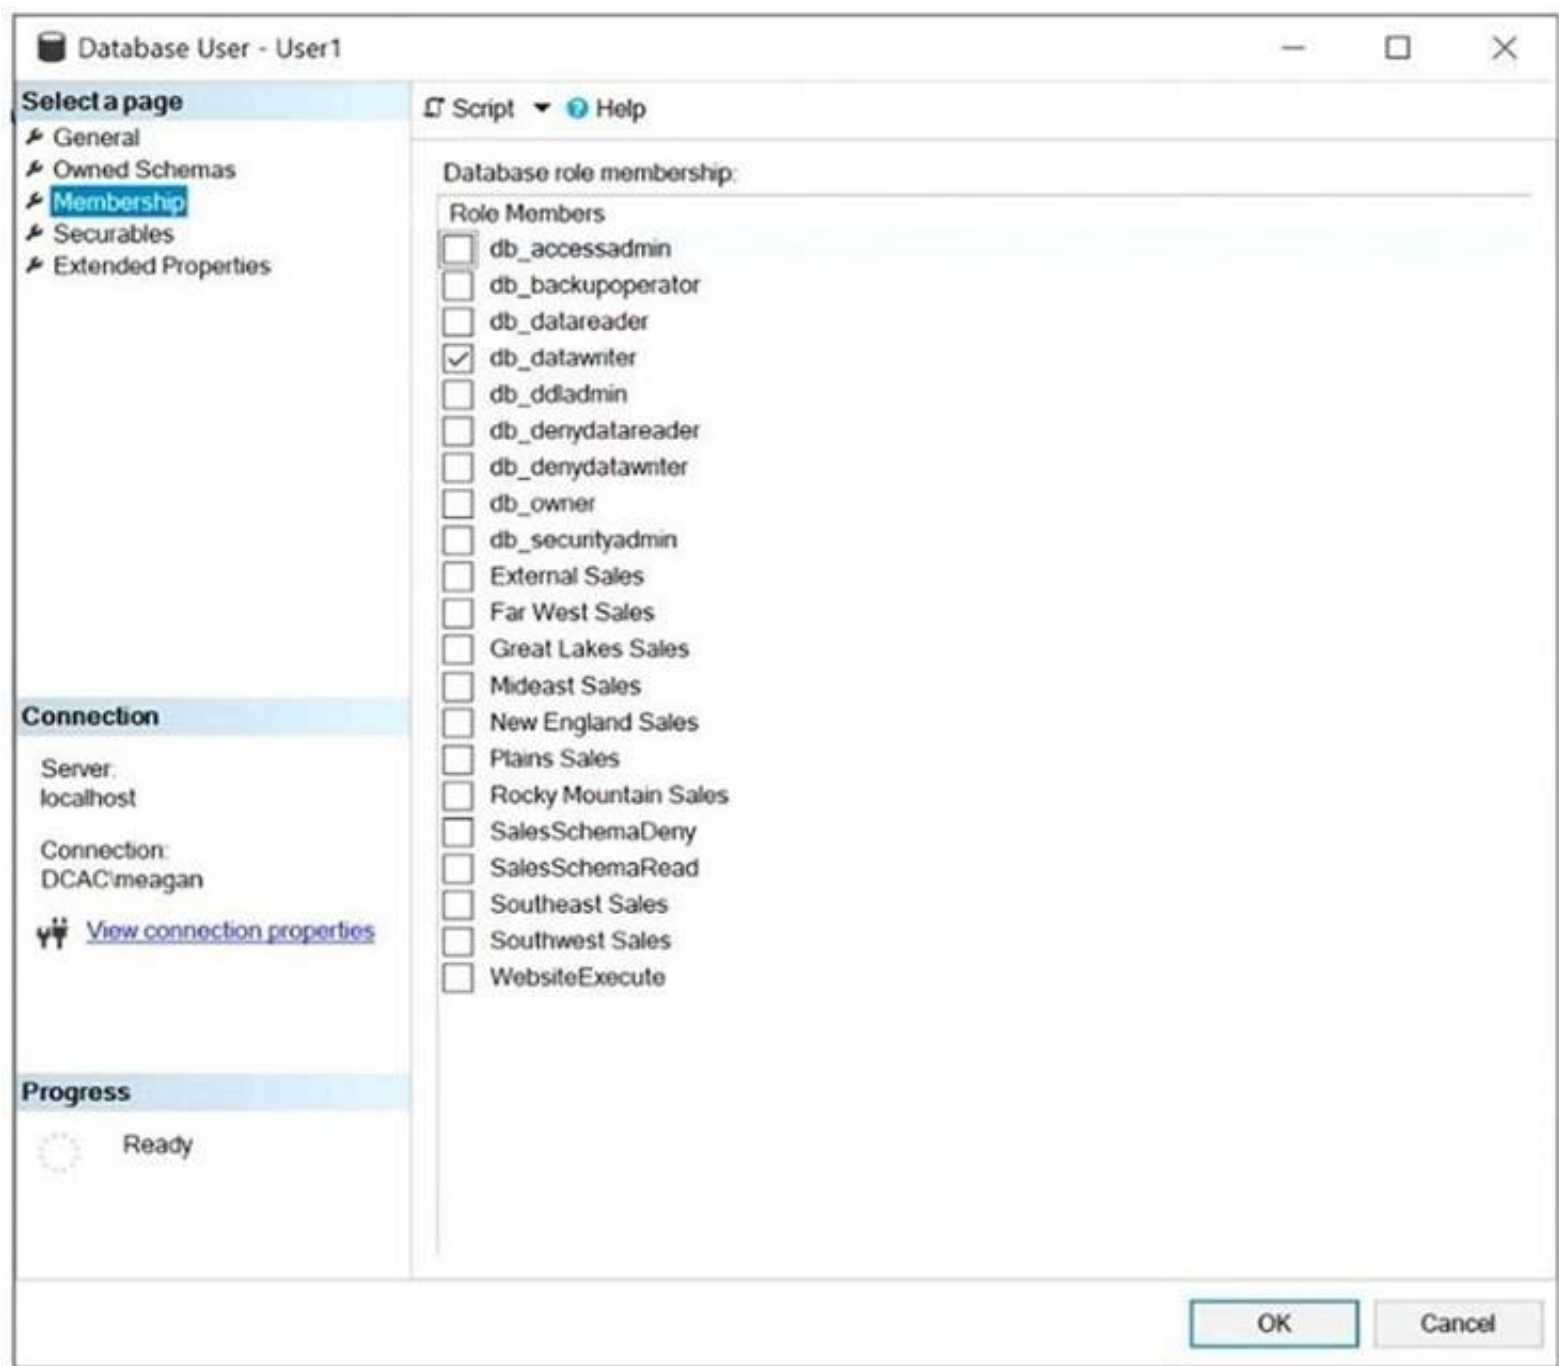
Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios>

**NEW QUESTION 57**

- (Exam Topic 5)

You have a Microsoft SQL Server database named DB1 that contains a table named Table1. The database role membership for a user named User1 is shown in the following exhibit.



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

User1 can [answer choice].

add a column to Table1  
delete a row from Table1  
delete Table1

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

db\_datareader  
db\_ddladmin  
db\_denydatareader  
db\_denydatawriter

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**  
Box 1: delete a row from Table1  
Members of the db\_datawriter fixed database role can add, delete, or change data in all user tables. Box 2: db\_datareader  
Members of the db\_datareader fixed database role can read all data from all user tables. Reference:  
<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

NEW QUESTION 61

- (Exam Topic 5)  
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.

**Answer:** AE

#### NEW QUESTION 64

- (Exam Topic 5)

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 Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes mapping data flow, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### Explanation:

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not a mapping flow, with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed. Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

#### NEW QUESTION 67

- (Exam Topic 5)

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

- The web app must be hosted on an Azure virtual network.
- The Azure SQL database must be assigned a private IP address.
- The Azure SQL database must allow connections only from the virtual network.

You need to recommend a solution that meets the requirements. What should you include in the recommendation?

- A. Azure Private Link
- B. a network security group (NSG)
- C. a database-level firewall
- D. a server-level firewall

**Answer:** A

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/private-endpoint-overview>

#### NEW QUESTION 70

- (Exam Topic 5)

You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. DF1 contains a linked service.

You have an Azure Key vault named vault1 that contains an encryption key named key1. You need to encrypt df1 by using key1.

What should you do first?

- A. Disable purge protection on vault1.



- B. Remove the linked service from df1.
- C. Create a self-hosted integration runtime.
- D. Disable soft delete on vault1.

**Answer:** B

**Explanation:**

A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.

Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key>

**NEW QUESTION 74**

- (Exam Topic 5)

You have the following Azure Resource Manager template.

```
...
"variable": {
  "serverName": "azsqlserver0001"
},
"resources": [
  {
    "name": "[variables('serverName')]",
    "type": "Microsoft.Sql/servers",
    "apiVersion": "2019-06-01-preview",
    "location": "[parameters('location')]",
    "properties": {
      "administratorLogin": "[parameters('administratorLogin')]",
      "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
      "version": "12.0"
    },
  },
  "resources": [
    {
      "name": "[concat(variables('serverName'), '/', parameters('databaseName'))]",
      "type": "Microsoft.Sql/servers/databases",
      "apiVersion": "2020-08-01-preview",
      "location": "[parameters('location')]",
      "kind": "v12.0"
      "sku": {
        "name": "Standard",
        "tier": "Standard",
        "capacity": 10
      },
      "dependsOn": [
        "[concat('Microsoft.Sql/servers/', variables('serverName'))]"
      ],
      "properties": {
      },
    },
    "resources": [
    ]
  }
]
},
],
...
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The template deploys a serverless Azure SQL database.	<input type="radio"/>	<input type="radio"/>
The template deploys a database to an Azure SQL Database managed instance.	<input type="radio"/>	<input type="radio"/>
The pricing tier of the database deployment is based on DTUs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

A screenshot of a computer Description automatically generated with low confidence

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/purchasing-models> <https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart>

NEW QUESTION 75

- (Exam Topic 5)

You have an Azure Databricks resource.

You need to log actions that relate to changes in compute for the Databricks resource. Which Databricks services should you log?

- A. clusters
- B. jobs
- C. DBFS
- D. SSH
- E. workspace

Answer: E

Explanation:

Cloud Provider Infrastructure Logs.

Databricks logging allows security and admin teams to demonstrate conformance to data governance standards within or from a Databricks workspace.

Customers, especially in the regulated industries, also need records on activities like:

- > User access control to cloud data storage
- > Cloud Identity and Access Management roles
- > User access to cloud network and compute

Azure Databricks offers three distinct workloads on several VM Instances tailored for your data analytics workflow—the Jobs Compute and Jobs Light Compute workloads make it easy for data engineers to build and execute jobs, and the All-Purpose Compute workload makes it easy for data scientists to explore, visualize, manipulate, and share data and insights interactively.

Reference:

<https://databricks.com/blog/2020/03/25/trust-but-verify-with-databricks.html>

NEW QUESTION 76

- (Exam Topic 5)

You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database. A recent failure causes the master database to be lost.

You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.

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

Actions

Answer Area

Add a certificate to an Azure key vault

Enable Transparent Data Encryption (TDE)

Encrypt a copy of the master key by using the service master key

Turn on the TRUSTWORTHY property and the CLR property

Attach the SSISDB database

Open the master key for the SSISDB database

>

<

↑

↓

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Attach the SSISDB database

Step 2: Turn on the TRUSTWORTHY property and the CLR property

If you are restoring the SSISDB database to an SQL Server instance where the SSISDB catalog was never created, enable common language runtime (clr)

Step 3: Open the master key for the SSISDB database

Restore the master key by this method if you have the original password that was used to create SSISDB. open master key decryption by password = 'LS1Setup!' --'Password used when creating SSISDB'

Alter Master Key Add encryption by Service Master Key

Step 4: Encrypt a copy of the mater key by using the service master key Reference:

<https://docs.microsoft.com/en-us/sql/integration-services/backup-restore-and-move-the-ssis-catalog>

NEW QUESTION 81

- (Exam Topic 5)

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 Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.  
 You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.  
 You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.  
 Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column.  
 Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:

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

**NEW QUESTION 82**

- (Exam Topic 5)

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.  
 Which resource provider should you enable?

- A. Microsoft.EventHub
- B. Microsoft.EventGrid
- C. Microsoft.Sql
- D. Microsoft.Automation

**Answer:** B

**Explanation:**

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger>

**NEW QUESTION 85**

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1.  
 You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.  
 How should you configure the backup of DB1? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Full and log backups only

Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn't support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.

Box 2: WITH CHECKSUM

Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database Migration Service only supports backups created using checksum.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online>



### NEW QUESTION 87

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool.

You run `PDW_SHOWSPACEUSED('dbo.FactInternetSales')`; and get the results shown in the following table.

ROWS	RESERVED_SPACE	DATA_SPACE	INDEX_SPACE	UNUSED_SPACE	PDW_NODE_ID	DISTRIBUTION_ID
694	2776	616	48	2112	1	1
407	2704	576	48	2080	1	2
53	2376	512	16	1848	1	3
58	2376	512	16	1848	1	4
168	2632	528	32	2072	1	5
195	2696	536	32	2128	1	6
5995	3464	1424	32	2008	1	7
0	2232	496	0	1736	1	8
264	2576	544	40	1992	1	9
3008	3016	960	32	2024	1	10
...	...	...	...	...	...	...
1550	2832	752	48	2032	1	50
1238	2832	696	40	2096	1	51
192	2632	528	32	2072	1	52
1127	2768	680	48	2040	1	53
1244	3032	704	64	2264	1	54
409	2632	568	32	2032	1	55
0	2232	496	0	1736	1	56
1437	2832	728	40	2064	1	57
0	2232	496	0	1736	1	58
384	2632	560	32	2040	1	59
225	2768	544	40	2184	1	60

Which statement accurately describes the `dbo.FactInternetSales` table?

- A. The table contains less than 10,000 rows.
- B. All distributions contain data.
- C. The table uses round-robin distribution
- D. The table is skewed.

**Answer: D**

#### Explanation:

The rows per distribution can vary up to 10% without a noticeable impact on performance. Here the distribution varies more than 10%. It is skewed.

Note: `SHOWSPACEUSED` displays the number of rows, disk space reserved, and disk space used for a specific table, or for all tables in a Azure Synapse Analytics or Parallel Data Warehouse database.

This is a very quick and simple way to see the number of table rows that are stored in each of the 60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.

`ROUND_ROBIN` distributed tables should not be skewed. Data is distributed evenly across the nodes by design.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu> <https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distrib>

### NEW QUESTION 90

- (Exam Topic 5)

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

**Answer: A**

#### Explanation:

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>

### NEW QUESTION 92

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL managed instance named `SQLMi1` and a SQL Agent job named `Backupdb`. `Backupdb` performs a daily backup of the databases hosted on `SQLMi1`.

You need to be notified by email if the job fails.

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.



Actions

Create a SQL Server Agent alert.

Create an operator.

Create an extended event.

Enable Database Mail.

Add a failure notification to the job.

Answer Area

<

>

↑

↓

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

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

NEW QUESTION 97

- (Exam Topic 5)

You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1. You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

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

NOTE: Each correct selection is worth one point.

- A. From the properties of DB1. configure geo-replication.
- B. From the properties of Server1 add a failover group.
- C. Create a new Azure SQL Database server named Server2.
- D. From the properties of Server1 configure retention for DB1
- E. Create a new Azure SQL Database instance named DB2.

Answer: BC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell> <https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azur>

NEW QUESTION 100

- (Exam Topic 5)

You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1.

You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.

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.

Actions

Run CREATE LOGIN user1@contoso.com FROM EXTERNAL PROVIDER on the master database.

Run ALTER SERVER ROLE securityadmin ADD MEMBER user1@contoso.com.

Create a managed identity for SQLMI1.

Grant SQLMI1 read access to Azure AD.

Run CREATE USER user1@contoso.com FROM LOGIN user1@contoso.com.

Answer Area

<

>

↑

↓

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated with medium confidence

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/aad-security-configure-tutorial>

**NEW QUESTION 105**

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

You need to add a 4-TB volume that meets the following requirements:

- Maximizes IOPs
- Uses premium solid state drives (SSDs)

What should you do?

- A. Attach two mirrored 4-TB SSDs.
- B. Attach a stripe set that contains four 1-TB SSDs.
- C. Attach a RAID-5 array that contains five 1-TB SSDs.
- D. Attach a single 4-TB SSD.

**Answer:** B

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/storage-configuration?tabs=window>

**NEW QUESTION 109**

- (Exam Topic 5)

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

- Support failovers that do not require client applications to change their connection strings.
- Replicate the database to a secondary Azure region.
- 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

**Answer:** A

**Explanation:**

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>

**NEW QUESTION 114**

- (Exam Topic 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

**Answer:** AC

**Explanation:**

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;.

Reference:

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

**NEW QUESTION 115**

- (Exam Topic 5)

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.

You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

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

correct order.

Actions

Deploy Grafana to an Azure virtual machine.

Build a `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file.

Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated with medium confidence  
Send application metrics using Dropwizard.  
Spark uses a configurable metrics system based on the Dropwizard Metrics Library.  
To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps: Step 1: Configure your Azure Databricks cluster to use the Databricksmonitoring library.  
Prerequisite: Configure your Azure Databricks cluster to use the monitoring library. Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file  
Step 3: Create Dropwizard counters in your application code Create Dropwizard gauges or counters in your application code

NEW QUESTION 119

- (Exam Topic 5)  
You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:

- > ProductID
- > ItemPrice
- > LineTotal
- > Quantity
- > StoreID
- > Minute
- > Month
- > Hour
- > Year
- > Day

You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.  
How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

df.write

.bucketBy

.partitionBy

.range

.sortBy

("")

("StoreID", "Hour")

("StoreID", "Year", "Month", "Day", "Hour")

("Year", "Month", "Day", "Hour", "StoreID")

.mode("append")

.csv("/Purchases")

.json("/Purchases")

.parquet("/Purchases")

.saveAsTable("/Purchases")

- A. Mastered

B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

Box 1: .partitionBy Example:

df.write.partitionBy("y","m","d") mode(SaveMode.Append) parquet("/data/hive/warehouse/db\_name.db/" + tableName) Box 2: ("Year","Month","Day","Hour","StoreID")

Box 3: .parquet("/Purchases") Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partiti>

**NEW QUESTION 123**

- (Exam Topic 5)

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.

You need to ensure that each customer can only connect to and access their respective database. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Implement row-level security (RLS).
- B. Create users in each database.
- C. Configure the database firewall.
- D. Configure the server firewall.
- E. Create logins in the master database.
- F. Implement Always Encrypted.

**Answer:** BC

**Explanation:**

Manage database access by adding users to the database, or allowing user access with secure connection strings.

Database-level firewall rules only apply to individual databases. Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/secure-database-tutorial>

**NEW QUESTION 125**

- (Exam Topic 5)

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 126**

- (Exam Topic 5)

You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

- A. An Azure SQL Database elastic pool
- B. SQL Server on Azure virtual machines
- C. an Azure SQL Database managed instance
- D. Azure SQL databases

**Answer:** A

**Explanation:**

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview>

**NEW QUESTION 127**

- (Exam Topic 5)

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 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 change the data file for the master database to autogrow by 10 percent. Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Reference:

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

**NEW QUESTION 132**

- (Exam Topic 5)

You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation.

Which service should you use to ingest the data?

- A. Azure Event Hubs Standard
- B. Azure Stream Analytics
- C. Azure Data Factory
- D. Azure Event Hubs Dedicated

**Answer: D**

**Explanation:**

The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.

Reference:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

**NEW QUESTION 133**

- (Exam Topic 5)

You are planning disaster recovery for the failover group of an Azure SQL Database managed instance.

Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs.

You set the Read/Write failover policy to Automatic.

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

NOTE: Each correct selection is worth one point.

- A. In the event of a datacenter or Azure regional outage, the databases will fail over automatically.
- B. In the event of an outage, the databases in the primary instance will fail over immediately.
- C. In the event of an outage, you can selectively fail over individual databases.
- D. In the event of an outage, you can set a different grace period to fail over each database.
- E. In the event of an outage, the minimum delay for the databases to fail over in the primary instance will be one hour.

**Answer: AE**

**Explanation:**

A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region.

E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

**NEW QUESTION 136**

- (Exam Topic 5)

You need to migrate an on-premises Microsoft SQL Server database to Azure SQL Database. The solution must minimize downtime.

What should you do?

- A. Configure Transaction Log Shipping.
- B. Implement Always On availability groups.
- C. Configure transactional replication.
- D. Import a BACPAC.

**Answer: C**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-to-database-from-sql-server#method-1-migra>

**NEW QUESTION 141**

- (Exam Topic 5)

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 two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: You run the Remove-AzSqlDatabase PowerShell cmdlet for Database1 on Server2. You run the Restore-AzSqlDatabase PowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

**NEW QUESTION 144**

- (Exam Topic 5)

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:

- During normal operations, provide at least two readable copies of Sales.
- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and failover groups. Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: 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>

**NEW QUESTION 145**

- (Exam Topic 5)

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.

You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.

What should you use?

- A. Performance Recommendations in the Azure portal
- B. Extended Events in Microsoft SQL Server Management Studio (SSMS)
- C. Query Store in Microsoft SQL Server Management Studio (SSMS)
- D. Query Performance Insight in the Azure portal

**Answer: C**

**Explanation:**

Use the Query Store Page in SQL Server Management Studio.

Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query.

This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu>

**NEW QUESTION 148**

- (Exam Topic 5)

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

- Send the output to an Azure Synapse.
- Identify spikes and dips in time series data.
- Minimize development and configuration effort. Which should you include in the solution?

- A. Azure SQL Database
- B. Azure Databricks
- C. Azure Stream Analytics

**Answer: C**

**Explanation:**

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>

#### NEW QUESTION 150

- (Exam Topic 5)

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

**Answer:** CD

#### Explanation:

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 non-zero, you should scale out your job, by increasing the SUs.

Reference:

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

#### NEW QUESTION 155

- (Exam Topic 5)

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 Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

#### Explanation:

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.

Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

#### NEW QUESTION 159

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains a nonclustered index named index1. End users report slow queries when they use index1.

You need to identify the operations that are being performed on the index. Which dynamic management view should you use?

- A. `sys.dm_exec_query_plan_stats`
- B. `sys.dm_db_index_physical_stats`
- C. `sys.dm_db_index_operational_stats`
- D. `sys.dm_db_index_usage_stats`

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** D

#### NEW QUESTION 160

- (Exam Topic 5)

You have an Azure SQL managed instance.  
You need to enable SQL Agent Job email notifications. What should you do?

- A. Use the Agent XPs option.
- B. Enable the SQL Server Agent.
- C. Run the sp\_configure command.
- D. Run the sp\_set\_agent\_properties command.

**Answer:** C

#### NEW QUESTION 165

- (Exam Topic 5)

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.

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	

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Graphical user interface, table Description automatically generated

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>

#### NEW QUESTION 168

- (Exam Topic 5)

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. \[<language>]
- B. %<language>
- C. \[<language>]
- D. @<language>

**Answer:** B

#### Explanation:

You can override the default language by specifying the language magic command %<language> 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>

#### NEW QUESTION 170

- (Exam Topic 5)

You have an Azure Data Factory that contains 10 pipelines.

You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.

What should you add to each pipeline?

- A. an annotation
- B. a resource tag
- C. a run group ID
- D. a user property
- E. a correlation ID

**Answer:** A

#### Explanation:

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.

Reference:

<https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/>

#### NEW QUESTION 172

- (Exam Topic 5)

You have a Microsoft SQL Server 2017 server.

You need to migrate the server to Azure. The solution must meet the following requirements:

- Ensure that the latest version of SQL Server is used.
- Support the SQL Server Agent service. Minimize administrative effort.

What should you use?

- A. SQL Server on Azure Virtual Machines
- B. Azure SQL Database
- C. an Azure SQL Database elastic pool
- D. Azure SQL Managed Instance

**Answer:** A

#### NEW QUESTION 175

- (Exam Topic 5)

You have a version-8.0 Azure Database for MySQL database.

You need to identify which database queries consume the most resources. Which tool should you use?

- A. Query Store
- B. Metrics
- C. Query Performance Insight
- D. Alerts

**Answer:** A

#### Explanation:

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time. Query Store simplifies performance troubleshooting by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns.

Data for all users, databases, and queries is stored in the mysql schema database in the Azure

Database for MySQL instance. Reference:

<https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store>

#### NEW QUESTION 178

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.

You need to create a data source in Pool1 that you can reference when you create the external table. 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.

CREATE EXTERNAL DATA SOURCE source1

WITH

( LOCATION = 'https://account1.

	▼
blob	
dfs	
table	

.core.windows.net',

	▼
PUSHDOWN = ON	
TYPE = BLOB_STORAGE	
TYPE = HADOOP	

)

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, table Description automatically generated

Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2 CREATE EXTERNAL DATA SOURCE YellowTaxi

WITH ( LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/', TYPE = HADOOP)

Box 2: HADOOP

Reference:

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

**NEW QUESTION 183**

- (Exam Topic 5)

You have the following Azure Data Factory pipelines:

- > Ingest Data from System1
- > Ingest Data from System2
- > Populate Dimensions
- > Populate Facts

Ingest Data from System1 and Ingest Data from System2 have no dependencies. Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2. Populate Facts must execute after the Populate Dimensions pipeline. All the pipelines must execute every eight hours.

What should you do to schedule the pipelines for execution?

- A. Add a schedule trigger to all four pipelines.
- B. Add an event trigger to all four pipelines.
- C. Create a parent pipeline that contains the four pipelines and use an event trigger.
- D. Create a parent pipeline that contains the four pipelines and use a schedule trigger.

**Answer:** D

**Explanation:**

Reference:

<https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/>

**NEW QUESTION 187**

- (Exam Topic 5)

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

**Answer:** C

**NEW QUESTION 189**

- (Exam Topic 5)

You have an Azure SQL Database managed instance. The instance starts experiencing performance issues.

You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.

What should you use?

- A. the Azure portal
- B. Extended Events
- C. Query Store
- D. dynamic management views

**Answer:** C

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu>

**NEW QUESTION 194**

- (Exam Topic 5) You have an Azure SQL database. You identify a long running query.

You need to identify which operation in the query is causing the performance issue.

What should you use to display the query execution plan in Microsoft SQL Server Management Studio (SSMS)?

- A. Live Query Statistics
- B. an estimated execution plan
- C. an actual execution plan
- D. Client Statistics

**Answer: C**

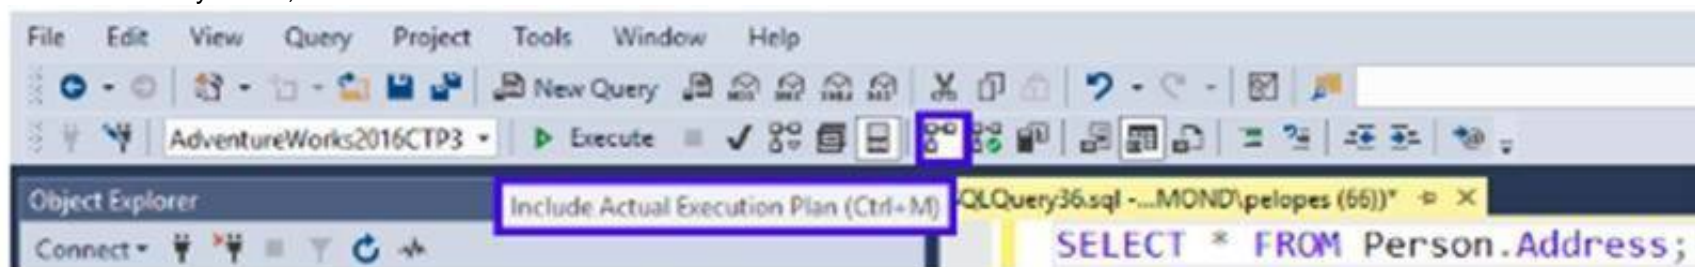
**Explanation:**

To include an execution plan for a query during execution

\* 1. On the SQL Server Management Studio toolbar, click Database Engine Query. You can also open an existing query and display the estimated execution plan by clicking the Open File toolbar button and locating the existing query.

\* 2. Enter the query for which you would like to display the actual execution plan.

\* 3. On the Query menu, click Include Actual Execution Plan or click the Include Actual Execution Plan toolbar button.



Note: Actual execution plans are generated after the Transact-SQL queries or batches execute. Because of this, an actual execution plan contains runtime information, such as actual resource usage metrics and runtime warnings (if any). The execution plan that is generated displays the actual query execution plan that the SQL Server Database Engine used to execute the queries.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/display-an-actual-execution-plan>

**NEW QUESTION 197**

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.

You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.

You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.

What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Run: A full backup and a log backup appended to the same file by using the WITH CHECKSUM option  
A full backup and a log backup to separate files by using the WITH CHECKSUM option  
A full backup and a log backup to separate files by using the WITH FILE\_SNAPSHOT option

Store the backup in: A Recovery Services vault  
An Azure Blob storage account  
An SMB file share

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Graphical user interface, application, Word Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-managed-instance-online>

**NEW QUESTION 202**

- (Exam Topic 5)

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.

#### Statements

#### Answer Area

```
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'
```

```
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
```

```
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'
```

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

```
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
```

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



- A. Mastered
- B. Not Mastered

Answer: A

#### Explanation:

Text Description automatically generated with low confidence

Text Description automatically generated

Reference:

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



### NEW QUESTION 206

- (Exam Topic 5)

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

➤ Restrict network connectivity to SqlSrv1.

What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private endpoint
- D. an ExpressRoute gateway

**Answer: C**

#### Explanation:

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.

The service could be an Azure service such as:

- Azure Storage
- Azure Cosmos DB
- Azure SQL Database
- Your own service using a Private Link Service. Reference:

<https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview>

### NEW QUESTION 208

- (Exam Topic 5)

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.

You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.

Which four 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.

#### Actions

#### Answer Area

Change the Active Directory Admin on TestServer1

Change the server name and related variables in the templates

From the database project, deploy the database schema and permissions

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

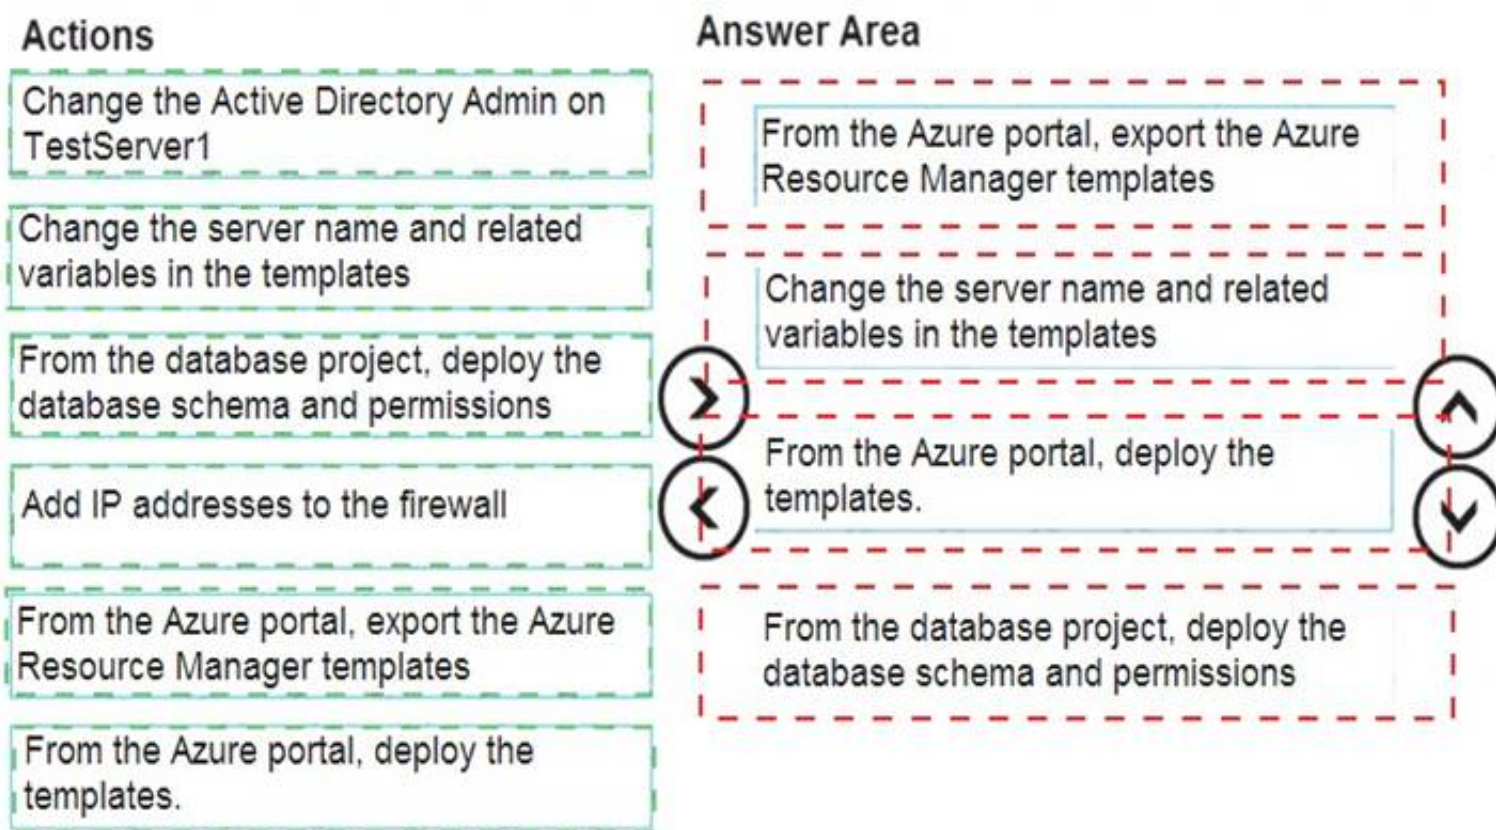
From the Azure portal, deploy the templates.



- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:



#### NEW QUESTION 210

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains a private certificate named Sales. The private key for Sales is encrypted with a password. You need to change the password for the private key. Which Transact-SQL statement should you run?

A)

```
ALTER CERTIFICATE Sales
WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'Mb^6BK&*w%',
ENCPTION BY PASSWORD = ' 6YY9YcD!pV');
```

B)

```
ALTER CERTIFICATE Sales
WITH PRIVATE KEY (ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');
```

C)

```
ALTER CERTIFICATE Sales WITH PRIVATE KEY (FILE = 'D:\importkeys\SalesNew, DECRYPTION BY PASSWORD = ' Mb^6BK&*w%');
```

D)

```
ALTER CERTIFICATE Sales WITH PRIVATE KEY (DECRYPTION BY PASSWORD = ' EWYx9Xk+$#');
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer: C**

#### NEW QUESTION 213

- (Exam Topic 5)

You create five Azure SQL Database instances on the same logical server.

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1. User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.

You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.

What should you do?

- A. Create User1 in the master database.
- B. Assign User1 the db\_datareader role for the master database.
- C. Assign User1 the db\_datareader role for the databases that User1 creates.
- D. Grant select on sys.databases to public in the master database.

**Answer: A**

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage>

#### NEW QUESTION 218

- (Exam Topic 5)

You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.

You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.



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.

### Actions

### Answer Area

Encrypt the Salary column by using the randomized encryption type.

Create a column encryption key.

Enable Transparent Data Encryption (TDE).

Encrypt the Salary column by using the deterministic encryption type.

Apply a dynamic data mask to the Salary column.

Create a column master key.



- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Step 1: Create a column master key

Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.

Step 2: Create a column encryption key.

Step 3: Encrypt the Salary column by using the randomized encryption type.

Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

#### NEW QUESTION 223

- (Exam Topic 5)

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:

- > During normal operations, provide at least two readable copies of Sales.
- > Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo-replication. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: 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>

#### NEW QUESTION 226

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Premium
- B. Azure SQL Database serverless
- C. Azure SQL Database managed instance Business Critical

D. Azure SQL Database Standard

Answer: A

Explanation:

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). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.

Reference:

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

NEW QUESTION 230

- (Exam Topic 5)

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:

Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits. Supports backfilling existing data in the table.

Which type of trigger should you use?

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 235

- (Exam Topic 5)

You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.

You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1. Which four 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.

Actions

Create a workload group.

Create a user-defined classifier function.

Modify Resource Governor.

Create a contained database user.

Create a resource pool.

Answer Area

⬅

➡

⬆

⬆

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text, table Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server> <https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-def>

NEW QUESTION 240

- (Exam Topic 5)

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 two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: From the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.



Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

**NEW QUESTION 245**

- (Exam Topic 5)

You have an Azure SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1.

You need to test impersonation of user1 in db1 by running a SELECT statement and returning to the original execution context.

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.

EXECUTE AS 

	▼
CALLER	
LOGIN	
OWNER	
USER	

 = 'user1@contoso.com'

GO

SELECT SUSER\_SNAME ()

	▼
REVERT	
REVOKE	
ROLLBACK	

GO

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15> <https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-ssname-transact-sql?view=sql-server-ver15>

**NEW QUESTION 248**

- (Exam Topic 5)

You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.

You need to implement alerts by using Azure Monitor. The solution must meet the following requirements: ➤ Minimize costs.

➤ Aggregate Intelligent Insights telemetry from each database. What should you do?

- A. From the Diagnostic settings of each database, select Send to Log Analytics.
- B. From the Diagnostic settings of each database, select Stream to an event hub.
- C. From the Diagnostic settings of SQLMI1. select Send to Log Analytics.
- D. From the Diagnostic settings of SQLMI1. select Stream to an event hub.

**Answer:** A

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-expo>

**NEW QUESTION 252**

- (Exam Topic 5)

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 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 reduce the use of table variables and temporary tables. Does this meet the goal?

- A. Yes
- B. No

Answer: A

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

NEW QUESTION 254

- (Exam Topic 5)  
You have an Azure subscription that contains the resources shown in the following table.

Name	Type
App1	Azure web app
db1	Azure SQL database in the serverless tier

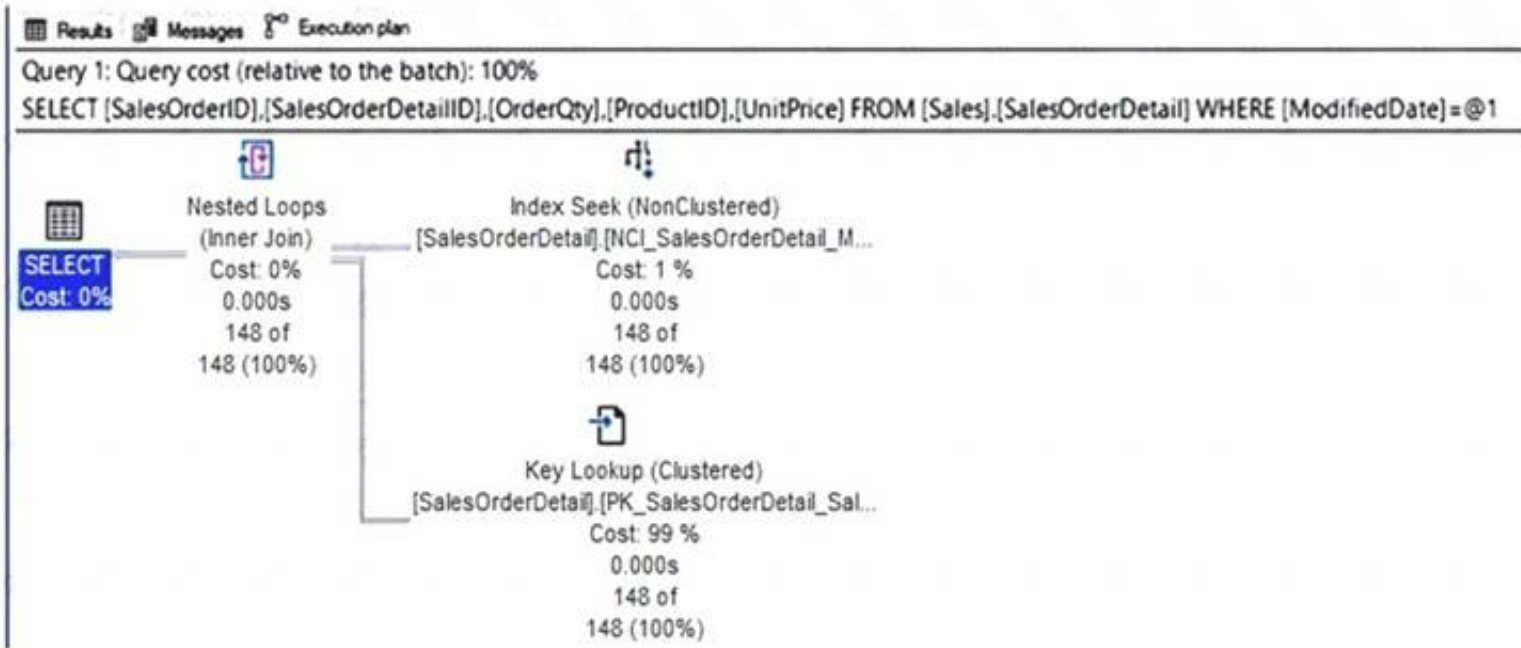
App1 experiences transient connection errors and timeouts when it attempts to access db1 after extended periods of inactivity. You need to modify db1 to resolve the issues experienced by App1 as soon as possible, without considering immediate costs. What should you do?

- A. Increase the number Of vCores allocated to db1.
- B. Disable auto-pause delay for db1.
- C. Decrease the auto-pause delay for db1.
- D. Enable automatic tuning for db1.

Answer: D

NEW QUESTION 255

- (Exam Topic 5)  
You have an Azure SQL database.  
You have a query and the associated execution plan as shown in the following exhibit.



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

The performance issue stems from the [answer choice] operator.

Select  
Index Seek  
Key Lookup  
Nested Loops

The performance issue can be resolved by adding include columns to the [answer choice].

heap  
clustered index  
nonclustered index

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

Box 1: Key Lookup

The Key Lookup cost is 99% so that is the performance bottleneck. Box 2: nonclustered index

The key lookup on the clustered index is used because the nonclustered index does not include the required columns to resolve the query. If you add the required columns to the nonclustered index, the key lookup will not be required.

NEW QUESTION 258

- (Exam Topic 5)

You have an Azure data factory that has two pipelines named PipelineA and PipelineB. PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



You create an alert for the data factory that uses Failed pipeline runs metrics for both pipelines and all failure types. The metric has the following settings:

- > Operator: Greater than
- > Aggregation type: Total
- > Threshold value: 2
- > Aggregation granularity (Period): 5 minutes
- > Frequency of evaluation: Every 5 minutes

Data Factory monitoring records the failures shown in the following table.

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input type="radio"/>

- A. Mastered  
B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Box 1: No

Just one failure within the 5-minute interval.

Box 2: No

Just two failures within the 5-minute interval.

Box 3: No

Just two failures within the 5-minute interval. Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview>

NEW QUESTION 260

- (Exam Topic 5)

You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1. You need to configure the autogrow and autoshrink settings for DB1.

Which statements should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.



Autogrow:

ALTER DATABASE MODIFY File and Filegroup options  
ALTER DATABASE SCOPED CONFIGURATION  
ALTER DATABASE SET options

Autoshrink:

ALTER DATABASE MODIFY File and Filegroup options  
ALTER DATABASE SCOPED CONFIGURATION  
ALTER DATABASE SET options

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

<https://learn.microsoft.com/en-us/troubleshoot/sql/admin/considerations-autogrow-autoshrink>

#### NEW QUESTION 262

- (Exam Topic 5)

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.

#### 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.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated

Reference:

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

#### NEW QUESTION 263

- (Exam Topic 5)

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

- Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
- Minimize the possibility of VM1 exfiltrating data stored in SqlDb1. What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link



D. an ExpressRoute gateway

**Answer: C**

**Explanation:**

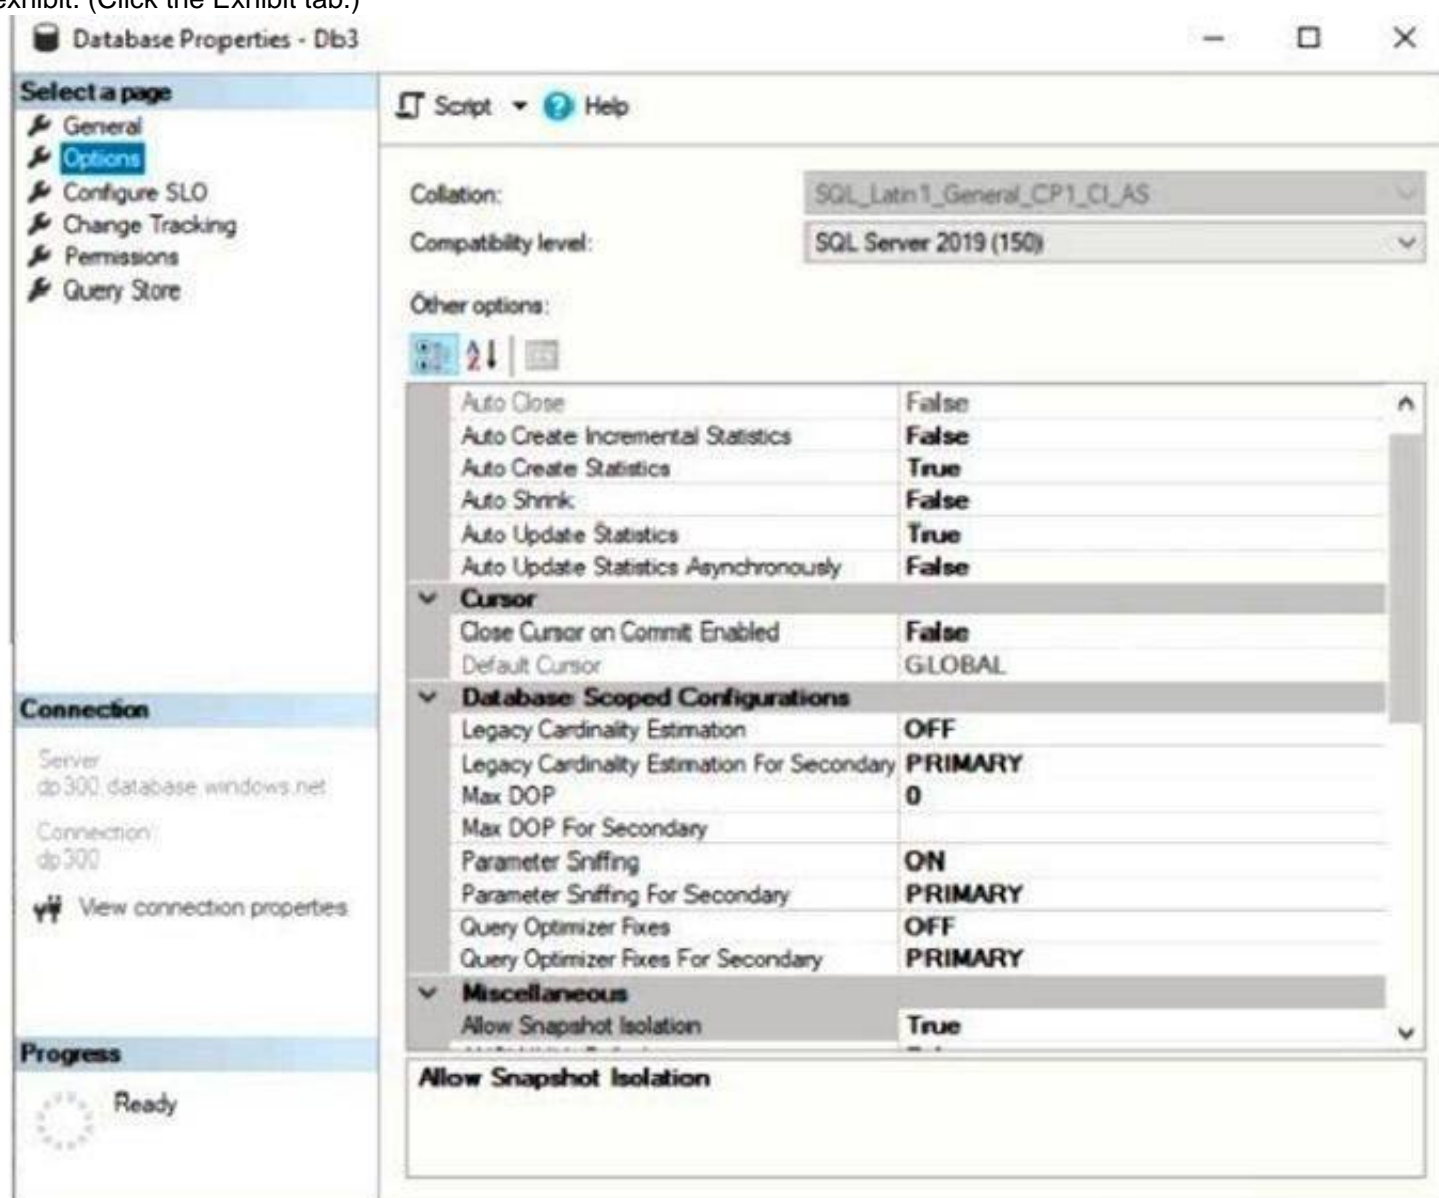
Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network. Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary. Reference: <https://docs.microsoft.com/en-us/azure/private-link/private-link-overview>

**NEW QUESTION 265**

- (Exam Topic 5)

You have an Azure SQL database named DB3.

You need to provide a user named DevUser with the ability to view the properties of DB3 from Microsoft SQL Server Management Studio (SSMS) as shown in the exhibit. (Click the Exhibit tab.)



Which Transact-SQL command should you run?

- A. GRANT SHOWPLAN TO DevUser
- B. GRANT VIEW DEFINITION TO DevUser
- C. GRANT VIEW DATABASE STATE TO DevUser
- D. GRANT SELECT TO DevUser

**Answer: C**

**Explanation:**

The exhibit displays Database [State] properties. To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission. Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-properties-options-page>

**NEW QUESTION 267**

- (Exam Topic 5)

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 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

**Answer: B**

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

**NEW QUESTION 268**

- (Exam Topic 5)  
You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. 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.

Actions

Duplicate Azure Policy definitions

Run Azure Policy remediation tasks

Create an Azure Blueprints assignment

Create an Azure Policy initiative

Create an Azure Policy initiative assignment

Answer Area

>

<

^

v

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**  
Step 1: Create an Azure Policy Initiative  
The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.  
With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.  
Step 2: Create an Azure Policy Initiative assignment  
Assign the initiative definition you created in the previous step. Step 3: Run Azure Policy remediation tasks  
To apply the Policy Initiative to the existing SQL databases. Reference:  
https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage

**NEW QUESTION 272**

- (Exam Topic 5)  
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 Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.  
You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.  
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.  
Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**  
Instead use a serverless SQL pool to create an external table with the extra column.  
Note: In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.  
Reference:  
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables

**NEW QUESTION 274**

- (Exam Topic 5)  
You have a SQL Server on Azure Virtual Machines instance named VM1 that hosts a database named DB1. You run the following query.

```
BACKUP LOG DB1 TO DISK = '\\File1\SQLBackups\DB1.trn'
WITH NORECOVERY,COPY_ONLY,CONTINUE_AFTER_ERROR;
GO
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

#### Answer Area

Statements	Yes	No
The log file will be truncated.	<input type="radio"/>	<input type="radio"/>
DB1 will be placed in an offline state.	<input type="radio"/>	<input type="radio"/>
You are performing a tail-log backup.	<input type="radio"/>	<input type="radio"/>

- A. Mastered  
B. Not Mastered

Answer: A

#### Explanation:

##### Answer Area

Statements	Yes	No
The log file will be truncated.	<input checked="" type="radio"/>	<input type="radio"/>
DB1 will be placed in an offline state.	<input type="radio"/>	<input checked="" type="radio"/>
You are performing a tail-log backup.	<input type="radio"/>	<input checked="" type="radio"/>

#### NEW QUESTION 279

- (Exam Topic 5)

You have an Always On availability group deployed to Azure virtual machines. The availability group contains a database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica.

You need to initiate a full backup of DB1 on SQL2. Which statement should you run?

- A. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (Differential, STATS=5, COMPRESSION);  
B. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (COPY\_ONLY, STATS=5, COMPRESSION);  
C. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (File\_Snapshot, STATS=5, COMPRESSION);  
D. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (NoInit, STATS=5, COMPRESSION);

Answer: B

#### Explanation:

BACKUP DATABASE supports only copy-only full backups of databases, files, or filegroups when it's executed on secondary replicas. Copy-only backups don't impact the log chain or clear the differential bitmap.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on>

#### NEW QUESTION 282

- (Exam Topic 5)

You have an Azure SQL database named sqldb1.

You need to minimize the amount of space by the data and log files of sqldb1. What should you run?

- A. DBCC SHRINKDATABASE  
B. sp\_clean\_db\_free\_space  
C. sp\_clean\_db\_file\_free\_space  
D. DBCC SHRINKFILE

Answer: A

#### Explanation:

DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database. Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql>

#### NEW QUESTION 287

- (Exam Topic 5)

You have an Azure Data Lake Storage Gen2 container.

Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that. Users can read files in the container but cannot modify the files.

You need to design a data archiving solution that meets the following requirements:

- New data is accessed frequently and must be available as quickly as possible.
- Data that is older than five years is accessed infrequently but must be available within one second when requested.
- Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.



➤ Costs must be minimized while maintaining the required availability.  
How should you manage the data? To answer, select the appropriate options in the answer area.  
NOTE: Each correct selection is worth one point.

Five-year-old data:

	▼
Delete the blob.	
Move to archive storage.	
Move to cool storage.	
Move to hot storage.	

Seven-year-old data:

	▼
Delete the blob.	
Move to archive storage.	
Move to cool storage.	
Move to hot storage.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text, table Description automatically generated

Box 1: Move to cool storage

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

Short-term backup and disaster recovery

Older data not used frequently but expected to be available immediately when accessed

Large data sets that need to be stored cost effectively, while more data is being gathered for future processing Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Move to archive storage

Example usage scenarios for the archive access tier include: Long-term backup, secondary backup, and archival datasets

Original (raw) data that must be preserved, even after it has been processed into final usable form Compliance and archival data that needs to be stored for a long time and is hardly ever accessed Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**NEW QUESTION 289**

- (Exam Topic 5)

You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL Server 2017 Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux. You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

To prepare for the migration:

	▼
--	---

To perform the migration, use:

	▼
--	---

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

To prepare for the migration:

Create a SQL Server 2019 Always On availability group on VM1.	▼
---	---

To perform the migration, use:

Azure Migrate	▼
---------------	---

**NEW QUESTION 292**

- (Exam Topic 5)



You have an Azure SQL managed instance.

You need to gather the last execution of a query plan and its runtime statistics. The solution must minimize the impact on currently running queries. What should you do?

- A. Generate an estimated execution plan.
- B. Generate an actual execution plan.
- C. Run sys.dm\_exec\_query\_plan\_scans.
- D. Generate Live Query Statistics.

**Answer:** C

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-quer>

#### NEW QUESTION 297

- (Exam Topic 4)

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements. Which Azure Storage functionality should you include in the solution?

- A. time-based retention
- B. change feed
- C. lifecycle management
- D. soft delete

**Answer:** C

**Explanation:**

The lifecycle management policy lets you:

Delete blobs, blob versions, and blob snapshots at the end of their lifecycles Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts>

#### NEW QUESTION 300

- (Exam Topic 2)

You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements. What should you do?

- A. Create an Azure Monitor alert rule that has a static threshold and assign the alert rule to an action group.
- B. Add a diagnostic setting that logs QueryStoreRuntimeStatistics and streams to an Azure event hub.
- C. Add a diagnostic setting that logs Timeouts and streams to an Azure event hub.
- D. Create an Azure Monitor alert rule that has a dynamic threshold and assign the alert rule to an action group.

**Answer:** D

**Explanation:**

Reference:

<https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/>

#### NEW QUESTION 301

- (Exam Topic 2)

What should you use to migrate the PostgreSQL database?

- A. Azure Data Box
- B. AzCopy
- C. Azure Database Migration Service
- D. Azure Site Recovery

**Answer:** C

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/dms/dms-overview>

#### NEW QUESTION 302

- (Exam Topic 2)

Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

- A. Business Critical 4-vCore
- B. Hyperscale
- C. General Purpose v-vCore
- D. Serverless

**Answer:** A

**Explanation:**

There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed. Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview>

#### NEW QUESTION 304

- (Exam Topic 1)

You are evaluating the business goals.

Which feature should you use to provide customers with the required level of access based on their service agreement?

- A. dynamic data masking
- B. Conditional Access in Azure
- C. service principals
- D. row-level security (RLS)

**Answer: D**

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver15>

#### NEW QUESTION 307

- (Exam Topic 1)

You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.

What should you include in the plan?

- A. Configure the Deleted databases settings for ResearchSrvOL
- B. Deploy and configure an Azure Backup server.
- C. Configure the Advanced Data Security settings for ResearchDBL
- D. Configure the Manage Backups settings for ResearchSrvOL

**Answer: D**

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/long-term-backup-retention-configure>

#### NEW QUESTION 311

- (Exam Topic 1)

You create all of the tables and views for ResearchDB1.

You need to implement security for ResearchDB1. The solution must meet the security and compliance requirements.

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.

Actions	Answer Area
Run the Always Encrypted wizard.	
Create an Azure Key Vault instance and generate a secret.	
Create an Azure Key Vault instance and configure an access policy.	⬅️ ⬆️
Create an Azure AD managed identity.	➡️ ⬇️
Register ResearchApp1 to Azure AD.	

- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:

Graphical user interface, text, application Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=az>

#### NEW QUESTION 313

- (Exam Topic 1)

You need to identify the cause of the performance issues on SalesSQLDb1.

Which two dynamic management views should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. sys.dm\_pdw\_nodes\_tran\_locks
- B. sys.dm\_exec\_compute\_node\_errors
- C. sys.dm\_exec\_requests

D. sys.dm\_cdc\_errors  
E. sys.dm\_pdw\_nodes\_os\_wait\_stats  
F. sys.dm\_tran\_locks

**Answer:** AE

**Explanation:**

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

A: Use sys.dm\_pdw\_nodes\_tran\_locks instead of sys.dm\_tran\_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information. SELECT

t1.resource\_type, t1.resource\_database\_id, t1.resource\_associated\_entity\_id, t1.request\_mode, t1.request\_session\_id, t2.blocking\_session\_id

FROM sys.dm\_tran\_locks as t1

INNER JOIN sys.dm\_os\_waiting\_tasks as t2

ON t1.lock\_owner\_address = t2.resource\_address;

Note: Depending on the system you're working with you can access these wait statistics from one of three locations:

sys.dm\_os\_wait\_stats: for SQL Server sys.dm\_db\_wait\_stats: for Azure SQL Database

sys.dm\_pdw\_nodes\_os\_wait\_stats: for Azure SQL Data Warehouse Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-lock>

**NEW QUESTION 317**

.....



## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DP-300 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DP-300 Product From:

<https://www.2passeasy.com/dumps/DP-300/>

## Money Back Guarantee

### DP-300 Practice Exam Features:

- \* DP-300 Questions and Answers Updated Frequently
- \* DP-300 Practice Questions Verified by Expert Senior Certified Staff
- \* DP-300 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DP-300 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year