

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 create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored You need to ensure that SQ Server Agent jobs running on SQL 1 can notify when a failure Occurs

Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.

Actions

Create a Database Mail account.

Enable pager notifications upon failure.

Create a profile named AzureManagedInstance_dbmail_profile.

Enable email notifications upon failure.

Create a profile named application_dbmail_profile.

Answer Area

➤
➤

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Create a Database Mail account.

Enable pager notifications upon failure.

Create a profile named AzureManagedInstance_dbmail_profile.

Enable email notifications upon failure.

Create a profile named application_dbmail_profile.

Answer Area

Create a Database Mail account.

Create a profile named AzureManagedInstance_dbmail_profile.

Enable email notifications upon failure.

➤
➤

NEW QUESTION 2

- (Exam Topic 5)

You have 20 Azure SQL databases provisioned by using the vCore purchasing model. You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Answer: ACE

Explanation:

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

For vCore-based purchasing model: MAX(<Total number of DBs X average vCore utilization per DB>,

<Number of concurrently peaking DBs X Peak vCore utilization per DB>)

A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

Reference:

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

NEW QUESTION 3

- (Exam Topic 5)

You have 10 Azure virtual machines that have SQL Server installed.

You need to implement a backup strategy to ensure that you can restore specific databases to other SQL Server instances. The solution must provide centralized management of the backups.

What should you include in the backup strategy?

- A. Automated Backup in the SQL virtual machine settings
- B. Azure Backup
- C. Azure Site Recovery
- D. SQL Server Agent jobs

Answer: B

Explanation:

Azure Backup provides an Enterprise class backup capability for SQL Server on Azure VMs. All backups are stored and managed in a Recovery Services vault.

There are several advantages that this solution provides, especially for Enterprises.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/backup-restore#azbackup>

NEW QUESTION 4

- (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 5

- (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 6

- (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 7

- (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 8

- (Exam Topic 5)

You have an Azure Synapse Analytics Apache Spark pool named Pool1.

You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.

You need to load the files into the tables. The solution must maintain the source data types. What should you do?

- A. Load the data by using PySpark.
- B. Load the data by using the OPENROWSET Transact-SQL command in an Azure Synapse Analytics serverless SQL pool.
- C. Use a Get Metadata activity in Azure Data Factory.
- D. Use a Conditional Split transformation in an Azure Synapse data flow.

Answer: B

Explanation:

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.

Serverless SQL pool enables you to query data in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.

To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.

The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files> <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-data-storage>

NEW QUESTION 9

- (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.

CREATE USER [App1] FROM login

[App1]
 [Contoso\app1]
 [App1@contoso.com]

login
 Windows
 EXTERNAL PROVIDER

- 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 10

- (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 10

- (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 14

- (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

CustomerKey

HASH

ROUND_ROBIN

REPLICATE

OrderDateKey

SalesOrderNumber

Answer Area

```
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 15

- (Exam Topic 5)
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

Actions

a virtual disk that uses the stripe layout

a virtual disk that uses the mirror layout

a volume

a virtual disk that uses the simple layout

a storage pool

Answer Area

>

<

↑

↓

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Follow these same steps to create striped virtual disk:

- Create Log Storage Pool.
- Create Virtual Disk
- Create Volume

Box 1: a storage pool

Box 2: a virtual disk that uses stripe layout

Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.

Box 3: a volume Reference:

<https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/>

NEW QUESTION 20

- (Exam Topic 5)

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

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

You need to resolve the performance issues. What should you do?

- A. Implement memory-optimized tables.
- B. Run the dbcc flushprocindb command.
- C. Replace the sequential index keys with nonsequential keys.
- D. Run the dbcc dbreindex command.

Answer: A

Explanation:

Reference:

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

NEW QUESTION 22

- (Exam Topic 5)

You have SQL Server on Azure virtual machines in an availability group. You have a database named DB1 that is NOT in the availability group.

You create a full database backup of DB1. You need to add DB1 to the availability group.

Which restore option should you use on the secondary replica?

- A. Restore with Recovery
- B. Restore with Norecovery
- C. Restore with Standby

Answer: B

Explanation:

Prepare a secondary database for an Always On availability group requires two steps:

* 1. Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY

* 2. Join the restored database to the availability group. Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server>

NEW QUESTION 26

- (Exam Topic 5)

You have an Azure virtual machine based on a custom image named VM1. VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements: Automate the patching of SQL Server and Windows Server.

Automate full database backups and transaction log backups of the databases on VM1.

Minimize administrative effort. What should you do first?

- A. Enable a system-assigned managed identity for VM1
- B. Register VM1 to the Microsoft.Sql resource provider
- C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1
- D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

Answer: B

Explanation:

Automated Patching depends on the SQL Server infrastructure as a service (IaaS) Agent Extension. The SQL Server IaaS Agent Extension (SqlIaaSExtension) runs on Azure virtual machines to automate administration

tasks. The SQL Server IaaS extension is installed when you register your SQL Server VM with the SQL Server VM resource provider.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extensionauto>

NEW QUESTION 30

- (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 31

- (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
```

CROSS APPLY
OPENJSON
OPENROWSET

- 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 32

- (Exam Topic 5)

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

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

▼ General

Operation Mode (Actual)Read only
Operation Mode (Requested)Read write

▼ Monitoring

Data Flush Interval (Minutes)15

Answer Area

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text Description automatically generated

NEW QUESTION 37

- (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 38

- (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 42

- (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 restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.
 Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

The REPLACE option overrides several important safety checks that restore normally performs. The overridden checks are as follows:

➤ Restoring over an existing database with a backup taken of another database.

With the REPLACE option, restore allows you to overwrite an existing database with whatever database is in the backup set, even if the specified database name differs from the database name recorded in the backup set. This can result in accidentally overwriting a database by a different database.

Reference:

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

NEW QUESTION 44

- (Exam Topic 5)

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state. What should you do?

- A. Double the value of Data Flush interval
- B. Decrease by half the value of Data Flush Interval
- C. Double the value of Statistics Collection Interval
- D. Decrease by half the value of Statistics Collection interval

Answer: B

Explanation:

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

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

NEW QUESTION 46

- (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 48

- (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 50

- (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 51

- (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 56

- (Exam Topic 5)

You have the following Azure Resource Manager template.

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 62

- (Exam Topic 5)

You have an Azure SQL managed instance named MI1.

You need to implement automatic tuning for the databases of MI1. What should you do?

- A. Use the REST API to call the patch operation and modify the AutomaticTuningServerMode property.
- B. Use Transact-SQL to enable the force_last_good_plan option.
- C. From the Azure portal, configure automatic tuning.

Answer: B

NEW QUESTION 66

- (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 70

- (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 75

- (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 76

- (Exam Topic 5)

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1. The only user who was created is the server administrator.

You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.

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
Connect to DB1 by using the Active Directory admin account.	
Create a user by using the FROM EXTERNAL PROVIDER clause.	
Connect to DB1 by using the server administrator account.	⬆
Set the Active Directory Admin for AzSQL1.	⬆
From the Azure portal, assign the SQL DB Contributor role to the user.	
Create a login in the master database.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Set up the Active Directory Admin for AzSQL1. Step 2: Connect to DB1 by using the server administrator.

Sign into your managed instance with an Azure AD login granted with the sysadmin role. Step 3: Create a user by using the FROM EXTERNAL PROVIDER clause.

FROM EXTERNAL PROVIDER is available for creating server-level Azure AD logins in SQL Database managed instance. Azure AD logins allow database-level Azure AD principals to be mapped to server-level Azure AD logins. To create an Azure AD user from an Azure AD login use the following syntax:

CREATE USER [AAD_principal] FROM LOGIN [Azure AD login] Reference:

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

NEW QUESTION 80

- (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 85

- (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	Answer Area
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.	

- 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 87

- (Exam Topic 5)

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

The Intelligent Insights diagnostics log identifies that several tables are missing indexes. You need to ensure that indexes are created for the tables. What should you do?

- A. Run the DBCC SQLPERF command.
 B. Run the dbcc dbreindex command.
 C. Modify the automatic tuning settings for db1.
 D. Modify the Query Store settings for db1.

Answer: C

Explanation:

Reference:

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

NEW QUESTION 90

- (Exam Topic 5)

You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.

ADF1 contains the following pipelines:

- P1: Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account
- P2: Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2

You need to configure P1 and P2 to maximize parallelism and performance.

Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

P1: ▼

- Set the Copy method to Bulk insert.
- Set the Copy method to PolyBase.
- Set the Isolation level to Repeatable read.
- Set the Partition option to Dynamic range.

P2: ▼

- Set the Copy method to Bulk insert.
- Set the Copy method to PolyBase.
- Set the Isolation level to Repeatable read.
- Set the Partition option to Dynamic range.

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

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

P1: Set the Partition option to Dynamic Range.

The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel. P2: Set the Copy method to PolyBase

Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store. (Polybase supports Azure Blob storage and Azure Data Lake Store by default.)

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse> <https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse>

NEW QUESTION 91

- (Exam Topic 5)

You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.

You plan to have a fact table for website visits. The table will be approximately 5 GB.

You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.

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

NOTE: Each correct selection is worth one point.

Distribution:

	▼
Hash	
Round robin	
Replicated	

Index:

	▼
Clustered columnstore	
Clustered	
Nonclustered	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: Hash

Consider using a hash-distributed table when:

The table size on disk is more than 2 GB.

The table has frequent insert, update, and delete operations. Box 2: Clustered columnstore

Clustered columnstore tables offer both the highest level of data compression and the best overall query performance.

Reference:

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

NEW QUESTION 93

- (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 96

- (Exam Topic 5)

You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

- A. automated
- B. interactive
- C. High Concurrency

Answer: A

Explanation:

Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.

The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/clusters>

NEW QUESTION 99

- (Exam Topic 5)

You are building a database backup solution for a SQL Server database hosted on an Azure virtual machine. In the event of an Azure regional outage, you need to be able to restore the database backups. The solution must minimize costs.

Which type of storage accounts should you use for the backups?

- A. locally-redundant storage (LRS)
- B. read-access geo-redundant storage (RA-GRS)
- C. zone-redundant storage (ZRS)
- D. geo-redundant storage

Answer: B

Explanation:

Geo-redundant storage (with GRS or GZRS) replicates your data to another physical location in the secondary region to protect against regional outages.

However, that data is available to be read only if the customer or Microsoft initiates a failover from the primary to secondary region. When you enable read access to the secondary region, your data is available to be read if the primary region becomes unavailable. For read access to the secondary region, enable read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Reference:

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

NEW QUESTION 100

- (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

Answer Area

- 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.



- 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 105

- (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 106

- (Exam Topic 5)

Your company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

A. Input Deserialization Errors

B. Late Input Events

C. Early Input Events

D. Watermark delay

Answer: D

Explanation:

The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.

The watermark delay metric can rise due to:

* 1. Not enough processing resources in Stream Analytics to handle the volume of input events.

* 2. Not enough throughput within the input event brokers, so they are throttled.

* 3. Output sinks are not provisioned with enough capacity, so they are throttled. Reference:

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

NEW QUESTION 107

- (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 108

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named Appl. Appl uses db1.

You need to enable Resource Governor for a App1. The solution must meet the following requirements: App1 must be able to consume all available CPU resources.

App1 must have at least half of the available CPU resources always available.

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	Answer Area
Create a plan.	
Create a classifier function in db1.	
Create a workload group.	
Create a classifier function in the master database.	
Create a resource pool that has the following configurations.	
<pre> MAX_CPU_PERCENT = 100 MIN_CPU_PERCENT = 50 </pre>	

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Create a plan.	Create a resource pool that has the following configurations.
Create a classifier function in db1.	<pre> MAX_CPU_PERCENT = 100 MIN_CPU_PERCENT = 50 </pre>
Create a workload group.	
Create a classifier function in the master database.	
Create a resource pool that has the following configurations.	
<pre> MAX_CPU_PERCENT = 100 MIN_CPU_PERCENT = 50 </pre>	

NEW QUESTION 113

- (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: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

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. Reference:

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

NEW QUESTION 115

- (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 117

- (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 122

- (Exam Topic 5)

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

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1.

The on-premises network connects to VNET1 by using an ExpressRoute connection.

You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service. You need to configure VNET1 to support the migration.

What should you do?

- A. Configure service endpoints.
- B. Configure virtual network peering.
- C. Deploy an Azure firewall.
- D. Configure network security groups (NSGs).

Answer: A

Explanation:

Reference:

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

NEW QUESTION 124

- (Exam Topic 5)

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10- Mbps internet connection.

You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database. What should you use?

- A. Azure Migrate
- B. Data Migration Assistant (DMA)
- C. Azure Data BOX
- D. Azure Database Migration Service

Answer: D

Explanation:

<https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Mig>

NEW QUESTION 127

- (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 132

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machine named SQL1.

You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort.

Where should you store the metrics?

- A. a Log Analytics workspace
- B. Azure Event Hubs
- C. Azure SQL Database
- D. an Azure Blob storage container

Answer: A

NEW QUESTION 133

- (Exam Topic 5)

You manage an enterprise data warehouse in Azure Synapse Analytics.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.

You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

- A. Local tempdb percentage
- B. DWU percentage
- C. Data Warehouse Units (DWU) used
- D. Cache hit percentage

Answer: A

Explanation:

Tempdb is used to hold intermediate results during query execution. High utilization of the tempdb database can lead to slow query performance.

Note: If you have a query that is consuming a large amount of memory or have received an error message related to allocation of tempdb, it could be due to a very large CREATE TABLE AS SELECT (CTAS) or INSERT SELECT statement running that is failing in the final data movement operation.

Reference:

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

NEW QUESTION 138

- (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 141

- (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 144

- (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 146

- (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 148

- (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 151

- (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 152

- (Exam Topic 5)

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

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

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

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

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

Answer: B

NEW QUESTION 156

- (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. two Azure SQL Databases in an elastic pool
- B. two databases hosted in SQL Server on an Azure virtual machine
- C. two databases in an Azure SQL Managed instance
- D. two single Azure SQL databases

Answer: D

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 161

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template.

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

NOTE: Each correct selection is worth one point.

```

"resources": [
  ...
  {
    "type": "Microsoft.Sql/servers",
    "name": "[parameters('targetName')]",
    "location": "[parameters('location')]",
    "sku": {
      "name": "[parameters('skuName')]"
    },
    "dependsOn": [
      "[parameters('targetName')]",
      "[parameters('virtualNetworkName')]",
      "[variables('networkSecurityGroupName')]"
    ],
    "properties": {
      "administratorLogin": "[parameters('administratorLogin')]",
      "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
      "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'), parameters('virtualNetworkName'), parameters('subnetName'))]",
      "storageSizeInGB": "[parameters('storageSizeInGB')]",
      "vCores": "[parameters('vCores')]",
      "licenseType": "[parameters('licenseType')]"
    }
  },
  ...
]

```

- <https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/create-template-quickstart?tabs=azure-powe>

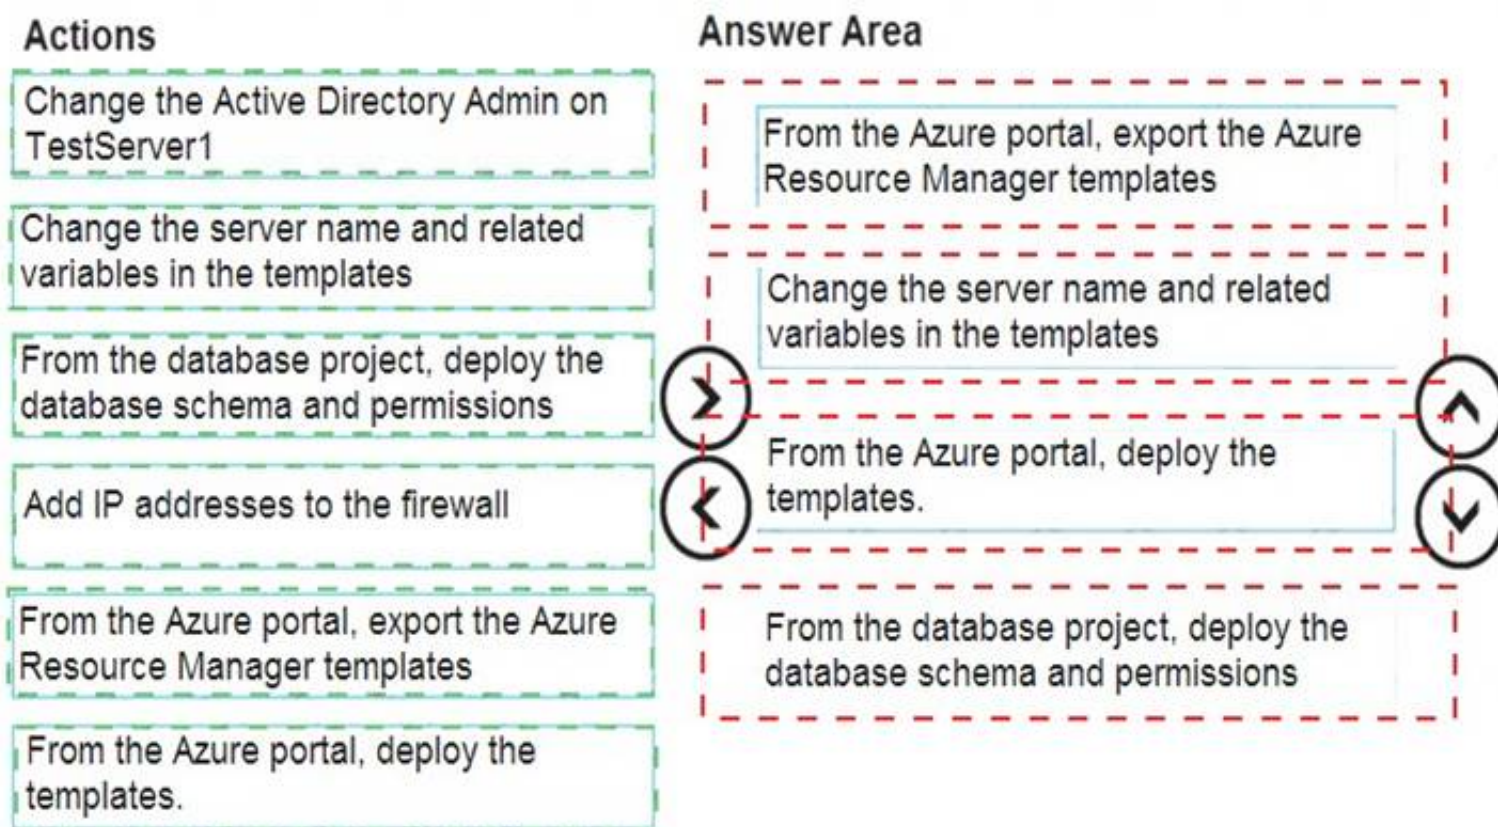
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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 165

- (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 166

- (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 170

- (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 172

- (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 173

- (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 178

- (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	Answer Area
Create a workload group.	
Create a user-defined classifier function.	
Modify Resource Governor.	
Create a contained database user.	
Create a resource pool.	

- 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 181

- (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 183

- (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-sname-transact-sql?view=sql-server-ver15>

NEW QUESTION 188

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named Db1.

You need to enable automatic tuning for Db1.

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

NOTE: Each correct selection is worth one point.

ALTER DATABASE [Db1]

	▼
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)	
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)	
SET AUTOMATIC_TUNING=AUTO	
SET QUERY_STORE=OFF	
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)	
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)	

GO

ALTER DATABASE [Db1]

	▼
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)	
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)	
SET AUTOMATIC_TUNING=AUTO	
SET QUERY_STORE=OFF	
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)	
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)	

GO

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SET AUTOMATIC_TUNING = AUTO

To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:

ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO

Setting automatic tuning to AUTO will apply Azure Defaults.

Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:

ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option. Setting it to OFF will also override any setting that database inherited and disable the tuning option.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable>

NEW QUESTION 193

- (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 194

- (Exam Topic 5)

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

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

What should you recommend?

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

Answer: C

Explanation:

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

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

Reference:

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

NEW QUESTION 197

- (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 199

- (Exam Topic 5)

You receive numerous alerts from Azure Monitor for an Azure SQL database.

You need to reduce the number of alerts. You must only receive alerts if there is a significant change in usage patterns for an extended period.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Set Threshold Sensitivity to High
B. Set the Alert logic threshold to Dynamic
C. Set the Alert logic threshold to Static
D. Set Threshold Sensitivity to Low
E. Set Force Plan to On

Answer: BD

Explanation:

B: Dynamic Thresholds continuously learns the data of the metric series and tries to model it using a set of algorithms and methods. It detects patterns in the data such as seasonality (Hourly / Daily / Weekly), and is able to handle noisy metrics (such as machine CPU or memory) as well as metrics with low dispersion (such as availability and error rate).

D: Alert threshold sensitivity is a high-level concept that controls the amount of deviation from metric behavior required to trigger an alert.

Low – The thresholds will be loose with more distance from metric series pattern. An alert rule will only trigger on large deviations, resulting in fewer alerts.

Reference:

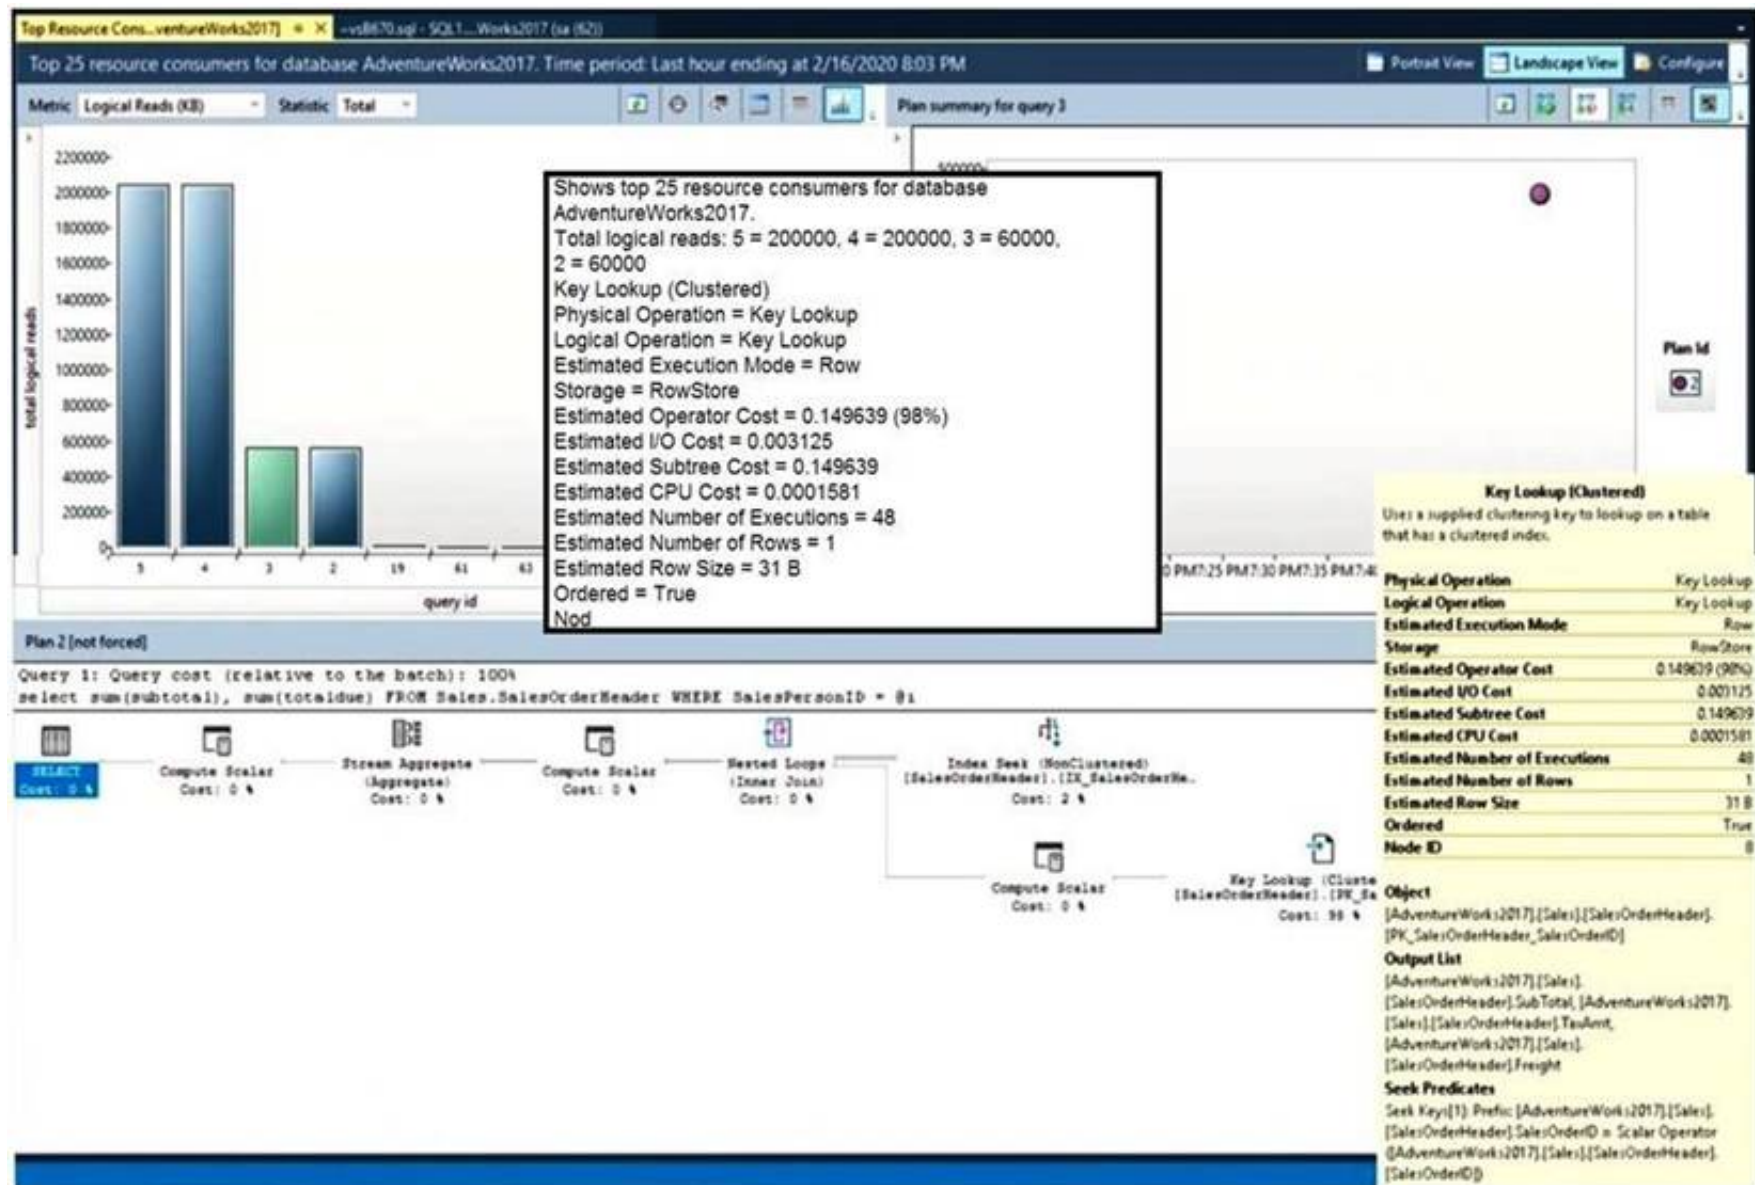
<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds>

NEW QUESTION 204

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



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

You will reduce the I/O usage and the query execution time if you force the query plan.

☐
☐

You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.

☐
☐

You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.

☐
☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Reference:

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

NEW QUESTION 207

- (Exam Topic 5)

You have an Azure SQL managed instance that hosts multiple databases.

You need to configure alerts for each database based on the diagnostics telemetry of the database. What should you use?

- A. Azure SQL Analytics alerts based on metrics
- B. SQL Health Check alerts based on diagnostics logs
- C. SQL Health Check alerts based on metrics
- D. Azure SQL Analytics alerts based on diagnostics logs

Answer: D

Explanation:

Reference:

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

NEW QUESTION 210

- (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 215

- (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 217

- (Exam Topic 5)

You have an Azure SQL database named db1.

You need to retrieve the resource usage of db1 from the last week.

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

NOTE: Each correct selection is worth one point.

SELECT *

FROM

	▼
sys.dm_db_resource_stats	
sys.dm_exec_requests	
sys.dm_user_db_resource_governance	
sys.resource_stats	

WHERE database_name = 'db1' AND

start_time >

	▼
DATEADD	
DATEDIFF	
DATEPART	
TOTETIMEOFFSET	

(day, -7, GETDATE())

ORDER BY start_time DESC;

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.resource_stats

sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.

Box 2: DateAdd

The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.

DECLARE @s datetime; DECLARE @e datetime;
SET @s= DateAdd(d,-7,GetUTCDate()); SET @e= GETUTCDATE();
SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats
WHERE start_time BETWEEN @s AND @e GROUP BY database_name
HAVING AVG(avg_cpu_percent) >= 80
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-data

NEW QUESTION 218

- (Exam Topic 5)
You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1.
You plan to create a database named DB1 in Pool1.
You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.
Which format should you use for the tables in DB1?

- A. JSON
- B. CSV
- C. Parquet
- D. ORC

Answer: C

Explanation:

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.
For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-storage-files-spark-tables

NEW QUESTION 222

- (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.

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:

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 227

- (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 232

- (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 235

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.

You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.

You need to ensure that webapp1 can identify the changes to Table1. The solution must meet the following requirements:

- Minimize compute times.
- Minimize storage.

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

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGETABLE()` function.

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON
```

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGE_TRACKING_CURRENT_VERSION()` function.

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_table
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_db
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING
```

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 239

- (Exam Topic 5)

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

You need to record whenever users query the CustomerPII table.

Which two options should you enable? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. server audit specification
- B. SQL Server audit
- C. database audit specification
- D. a server principal

Answer: AC

Explanation:

An auditing policy can be defined for a specific database or as a default server policy in Azure (which hosts SQL Database or Azure Synapse):

- A server policy applies to all existing and newly created databases on the server.
- If server auditing is enabled, it always applies to the database. The database will be audited, regardless of the database auditing settings.
- Enabling auditing on the database, in addition to enabling it on the server, does not override or change any of the settings of the server auditing. Both audits will exist side by side.

Note:

The Server Audit Specification object belongs to an audit.

A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created. Reference:

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

NEW QUESTION 240

- (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_scacs`.
- 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 244

- (Exam Topic 5)

You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies.

You need to ensure that users from each company can view only the data of their respective company. Which two objects should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. a column encryption key

- B. asymmetric keys
- C. a function
- D. a custom role-based access control (RBAC) role
- E. a security policy

Answer: DE

Explanation:

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes. Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac> <https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline>

NEW QUESTION 248

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1. Which isolation level should you configure?

- A. SERIALIZABLE
- B. SNAPSHOT
- C. READ UNCOMMITTED
- D. READ COMMITTED

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-data-sql-server-sql-database>

NEW QUESTION 249

- (Exam Topic 4)

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements. What should you create?

- A. a table that has a FOREIGN KEY constraint
- B. a table the has an IDENTITY property
- C. a user-defined SEQUENCE object
- D. a system-versioned temporal table

Answer: B

Explanation:

Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

Reference:

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

NEW QUESTION 253

- (Exam Topic 4)

You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.

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

NOTE: Each correct selection is worth one point.

Table type to store retail store data:

	▼
Hash	
Replicated	
Round-robin	

Table type to store promotional data:

	▼
Hash	
Replicated	
Round-robin	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: Hash Scenario:

Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.

A hash distributed table can deliver the highest query performance for joins and aggregations on large tables. Box 2: Round-robin Scenario:

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads. These are some scenarios where you should choose Round robin distribution:

- > When you cannot identify a single key to distribute your data.
- > If your data doesn't frequently join with data from other tables.
- > When there are no obvious keys to join.

Reference:

<https://rajanieshkaushikk.com/2020/09/09/how-to-choose-right-data-distribution-strategy-for-azure-synapse/>

NEW QUESTION 257

- (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 258

- (Exam Topic 2)

You are evaluating the role assignments.

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
DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.	<input type="radio"/>	<input type="radio"/>
DBAGroup1 will be able to assign the SQL DB Contributor role to other users.	<input type="radio"/>	<input type="radio"/>
DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

DBAGroup1 is member of the Contributor role.

The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.

Box 2: No

Box 3: Yes

DBAGroup2 is member of the SQL DB Contributor role.

The SQL DB Contributor role lets you manage SQL databases, but not access to them. Also, you can't manage their security-related policies or their parent SQL servers. As a member of this role you can create and manage SQL databases.

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

NEW QUESTION 259

- (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 263

- (Exam Topic 2)

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

- A. Azure Storage
- B. Azure Event Hubs
- C. Azure Log Analytics

Answer: C

Explanation:

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

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

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

Reference:

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

NEW QUESTION 264

- (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 267

- (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 270

- (Exam Topic 1)

You need to implement statistics maintenance for SalesSQLDb1. The solution must meet the technical requirements.

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 and configure a schedule.
- Create a SQL Server Agent job.
- Publish the runbook.
- Create an Azure Automation account.
- Import the SqlServer module.
- Create a runbook that runs a PowerShell script.
- Run `sp_add_jobserver`.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Automating Azure SQL DB index and statistics maintenance using Azure Automation:

- * 1. Create Azure automation account (Step 1)
- * 2. Import SQLServer module (Step 2)
- * 3. Add Credentials to access SQL DB

This will use secure way to hold login name and password that will be used to access Azure SQL DB

- * 4. Add a runbook to run the maintenance (Step 3)

Steps: * 1. Click on "runbooks" at the left panel and then click "add a runbook"

- * 2. Choose "create a new runbook" and then give it a name and choose "PowerShell" as the type of the runbook and then click on "create"

Add Runbook	Runbook
Quick Create Create a new runbook	<p>* Name ⓘ <input type="text" value="SqlMaintenance"/> ✓</p> <p>* Runbook type ⓘ <input type="text" value="PowerShell"/> ▼</p> <p>Description <input type="text" value=""/> ✓</p>
Import Import an existing runbook	

- * 5. Schedule task (Step 4)

Steps: 1. Click on Schedules 2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.

Reference:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statist>

NEW QUESTION 271

- (Exam Topic 1)

You need to configure user authentication for the SERVER1 databases. 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

Create a user in the master database

Modify the Azure SQL server administrator account

Create contained database users

Create an Azure AD administrator for the logical server

Connect to the databases by using an Azure AD account

Enable the contained database authentication option

Answer Area

>

<

↑

↓

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.
The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

- > Create and populate Azure AD.
- > Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
- > Create an Azure Active Directory administrator. (Step 1)
- > Configure your client computers.
- > Create contained database users in your database mapped to Azure AD identities. (Step 2)
- > Connect to your database by using Azure AD identities. (Step 3)

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

NEW QUESTION 272

- (Exam Topic 1)
You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Purchasing model:

Azure virtual machine reserved instances

DTU

vCore

Deployment option:

An Azure SQL Database elastic pool

An Azure SQL Database managed instance

A SQL Server Always On availability group

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: DTU
Scenario:

- > The 30 new databases must scale automatically.
- > Once all requirements are met, minimize costs whenever possible.

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.
In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.

Box 2: An Azure SQL database elastic pool
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. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

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

NEW QUESTION 276

- (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 277

.....

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