



Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (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: $\text{MAX}(\langle \text{Total number of DBs} \times \text{average vCore utilization per DB} \rangle, \langle \text{Number of concurrently peaking DBs} \times \text{Peak vCore utilization per DB} \rangle)$

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 2

- (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 3

- (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 4

- (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 5

- (Exam Topic 5)

You have an Azure SQL database named DB 1 in the General Purpose service tier. You need to monitor DB 1 by using SQL Insights.

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

To collect monitoring data, use:

To store monitoring data, create:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1 = Azure Monitor Agent Box 2 = An Azure SQL database

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql>

NEW QUESTION 6

- (Exam Topic 5)

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

Type	Designated retention period
Application	360 days
Infrastructure	60 days

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

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

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

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

NOTE: Each correct selection is worth one point.

To minimize storage costs:

Store the infrastructure logs and the application logs in the Archive access tier.

Store the infrastructure logs and the application logs in the Cool access tier.

Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

Azure Data Factory pipelines

Azure Blob storage lifecycle management rules

Immutable Azure Blob storage time-based retention policies

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A picture containing text Description automatically generated

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

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

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

Box 2: Azure Blob storage lifecycle management rules

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

Reference:

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

NEW QUESTION 7

- (Exam Topic 5)

You have an Azure SQL database.

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

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

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

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

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

Answer: C

Explanation:

OPTIMIZE_FOR_AD_HOC_WORKLOADS = { ON | OFF }

Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE_FOR_AD_HOC_WORKLOADS is enabled for a database, a compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.

Reference:

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

NEW QUESTION 8

- (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 9

- (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 10

- (Exam Topic 5)

You have the following Transact-SQL query.

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

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

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

Answer: C

Explanation:

Example:

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

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

Reference:

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

NEW QUESTION 10

- (Exam Topic 5)

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

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

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

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

NOTE: Each correct selection is worth one point.

Values

Answer Area

- CustomerKey
- HASH
- ROUND_ROBIN
- REPLICATE
- OrderDateKey
- SalesOrderNumber

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: HASH

Box 2: OrderDateKey

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

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

Reference:

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

NEW QUESTION 12

- (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	Answer Area
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	

- 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 14

- (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 16

- (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',

CROSS APPLY
 OPENJSON
 OPENROWSET

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: CREATE EXTERNAL TABLE

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

Syntax:

CREATE EXTERNAL TABLE { database_name.schema_name.table_name | schema_name.table_name | table_name }

(<column_definition> [,...n]) WITH (

LOCATION = 'folder_or_filepath', DATA_SOURCE = external_data_source_name, FILE_FORMAT = external_file_format_name

Box 2: OPENROWSET

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

Example: AS

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

OPENROWSET(BULK

'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*/

FORMAT='PARQUET') AS [r]

GROUP BY decennialTime, stateName GO

Reference:

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

NEW QUESTION 19

- (Exam Topic 5)

You have an Azure SQL Database instance named DatabaseA on a server named Server1.

You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.

You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.

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

On the master database, run CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;

On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';

On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];

On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';

On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'

Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:

```
CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
```

You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command.

Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:

```
CREATE USER readonlyuser FROM LOGIN readonlylogin;
```

Step 3: On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]

Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to the database via the db_datareader role.

```
EXEC sp_addrolemember 'db_datareader', 'readonlyuser';
```

Reference: <https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

NEW QUESTION 22

- (Exam Topic 5)

You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database. Which database option should set?

- A. PARAMETERIZATION to FORCED
- B. PARAMETERIZATION to SIMPLE
- C. Delayed Durability to Forced
- D. READ_COMMITTED_SNAPSHOT to ON

Answer: D

Explanation:

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

- > The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.
- > The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Reference:

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

NEW QUESTION 24

- (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 26

- (Exam Topic 5)

You have an Azure subscription.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

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

NEW QUESTION 29

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order.

You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders.

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.

Create:

- A materialized view in DB1
- A security policy in the Orders table**
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule**
- A table-valued function
- The CONTAINS predicate

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Create:

- A materialized view in DB1
- A security policy in the Orders table**
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule**
- A table-valued function**
- The CONTAINS predicate

NEW QUESTION 34

- (Exam Topic 5)

You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1. You need to ensure that an automatic email notification is sent once the job completes. What should you include in the solution?

- A. From SQL Server Configuration Manager (SSMS), enable SQL Server Agent
- B. From SQL Server Management Studio (SSMS), run sp_set_sqlagent_properties
- C. From SQL Server Management Studio (SSMS), create a Database Mail profile
- D. From the Azure portal, create an Azure Monitor action group that has an Email/SMS/Push/Voice action

Answer: C

Explanation:

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail. Using SQL Server Management Studio; to configure SQL Server Agent to use Database Mail:

- In Object Explorer, expand a SQL Server instance.
- Right-click SQL Server Agent, and then click Properties.
- Click Alert System.
- Select Enable Mail Profile.
- In the Mail system list, select Database Mail.
- In the Mail profile list, select a mail profile for Database Mail.
- Restart SQL Server Agent.

Note: Prerequisites include:

- Enable Database Mail.
- Create a Database Mail account for the SQL Server Agent service account to use.
- Create a Database Mail profile for the SQL Server Agent service account to use and add the user to the DatabaseMailUserRole in the msdb database.
- Set the profile as the default profile for the msdb database. Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d>

NEW QUESTION 39

- (Exam Topic 5)

You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases. You need to update the column and index statistics for the databases.

What should you use?

- A. an Azure Automation runbook
- B. a SQL Agent job
- C. Azure SQL Analytics
- D. automatic tuning in Azure SQL Database

Answer: A

Explanation:


Reference:
<https://www.sqlshack.com/automate-azure-sql-database-indexes-and-statistics-maintenance/>

NEW QUESTION 43

- (Exam Topic 5)

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

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



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

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

Answer: C

Explanation:

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

NEW QUESTION 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)

You have an Azure SQL database named db1 on a server named server1. You need to modify the MAXDOP settings for db1. What should you do?

- A. Connect to db1 and run the sp_configure command.
- B. Connect to the master database of server1 and run the sp_configure command.
- C. Configure the extended properties of db1.
- D. Modify the database scoped configuration of db1.

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/configure-max-degree-of-parallelism>

NEW QUESTION 53

- (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 54

- (Exam Topic 5)

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1. You need to encrypt the data in Col1. Which four actions should you perform for DB1 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 database master key.		
Create a column master key.		
Open the symmetric key.		
Create a certificate.		
Update Col1.		
Create a symmetric key.		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

Reference:

<https://www.sqlshack.com/an-overview-of-the-column-level-sql-server-encryption/>

NEW QUESTION 55

- (Exam Topic 5)

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory. You scale up the virtual machine to 8 vCPUs and 64 GB of memory. You need to provide the lowest latency for tempdb. What is the total number of data files that tempdb should contain?

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

Answer: C

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

Reference:

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

NEW QUESTION 56

- (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 57

- (Exam Topic 5)

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

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)

FactSales has 6 billion rows and is loaded nightly by using a batch process.

Which type of compression provides the greatest space reduction for the database?

- A. page compression
- B. row compression
- C. columnstore compression
- D. columnstore archival compression

Answer: D

Explanation:

Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.

Note: Columnstore — The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.

Reference:

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

NEW QUESTION 62

- (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 66

- (Exam Topic 5)

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.

Which resource provider should you enable?

A. Microsoft.EventHub

B. Microsoft.EventGrid

C. Microsoft.Sql

D. Microsoft.Automation

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 69

- (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 71

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool.

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

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

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

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

Answer: D

Explanation:

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

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

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

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

Reference:

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

NEW QUESTION 76

- (Exam Topic 5)

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

You need to create a surrogate key for the table. The solution must provide the fastest query performance. What should you use for the surrogate key?

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 77

- (Exam Topic 5)

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

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

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

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions	Answer Area
Create a SQL Server Agent alert.	
Create an operator.	⬅️ ⬆️
Create an extended event.	➡️ ⬇️
Enable Database Mail.	
Add a failure notification to the job.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

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

NEW QUESTION 81

- (Exam Topic 5)

You manage 100 Azure SQL managed instances located across 10 Azure regions.

You need to receive voice message notifications when a maintenance event affects any of the 10 regions. The solution must minimize administrative effort.

What should you do?

- A. From the Azure portal, create a service health alert.
- B. From the Azure portal, create an Azure Advisor operational excellence alert.
- C. From Microsoft SQL Server Management Studio (SSMS), configure a SQL Server agent job.
- D. From the Azure portal, configure an activity log alert.

Answer: C

NEW QUESTION 83

- (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 88

- (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 90

- (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 dbrindex 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 92

- (Exam Topic 5)

You have an Azure SQL database.

You run the following PowerShell script.

```
$serverName = "SERVER1"
$resourceGroup = "RG1"
$dbName = "DB1"

Connect-AzAccount

$server = Get-AzSqlServer -ServerName $serverName -ResourceGroupName
$resourceGroup

Set-AzSqlDatabaseBackupShortTermRetentionPolicy -ResourceGroupName $resourceGroup
-ServerName $server `
    -DatabaseName $dbName -RetentionDays 21

Set-AzSqlDatabaseBackupLongTermRetentionPolicy -ServerName $serverName -
DatabaseName $dbName `
    -ResourceGroupName $resourceGroup -WeeklyRetention P52W -YearlyRetention PSY
-WeekOfYear 52
```

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
DB1 can be restored to a specific point in time 30 days ago.	<input type="radio"/>	<input type="radio"/>
DB1 can be restored from a weekly backup performed six months ago.	<input type="radio"/>	<input type="radio"/>
DB1 can be restored from a yearly backup performed six years ago.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackupshorttermretentionpolicy?vi>

<https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackuplongtermretentionpolicy?vie>

NEW QUESTION 94

- (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 97

- (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 98

- (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 plan to migrate on-premises Microsoft SQL Server databases to Azure.

You need to identify which deployment and resiliency options meet the following requirements:

- > Support user-initiated backups.
- > Support multiple automatically replicated instances across Azure regions.
- > Minimize administrative effort to implement and maintain business continuity. What should you identify? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Deployment option:

	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SQL Server on Azure VMs

SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.

Box 2: Active geo-replication

Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-i> <https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>

NEW QUESTION 102

- (Exam Topic 5)

You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently.

You enable columnstore archival compression.

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

NOTE: Each correct selection is worth one point.

- A. Queries that use the index will consume more disk I/O.
- B. Queries that use the index will retrieve fewer data pages.
- C. The index will consume more disk space.
- D. The index will consume more memory.
- E. Queries that use the index will consume more CPU resources.

Answer: BE

Explanation:

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. In addition to saving space, data compression can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk.

Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data.

NEW QUESTION 105

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

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
Export a built-in policy.	
Create a custom policy based on a condition.	
Create a custom condition based on a built-in facet.	⬅️
View the policy history.	➡️
Import a policy file.	
Run a policy evaluation.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

<https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-bas>

NEW QUESTION 107

- (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 110

- (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 113

- (Exam Topic 5)

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

Which service should you use to ingest the data?

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

Answer: D

Explanation:

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

Reference:

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

NEW QUESTION 114

- (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 118

- (Exam Topic 5)

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

What should you do?

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 119

- (Exam Topic 5)

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

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.

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

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

> Ensure that Sales remains available if a datacenter fails.

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

- A. Yes
- B. No

Answer: B

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

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

Reference:

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

NEW QUESTION 120

- (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 125

- (Exam Topic 5)

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

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

A. `sys.dm_exec_query_plan_stats`

B. `sys.dm_db_index_physical_stats`

C. `sys.dm_db_index_operational_stats`

D. `sys.dm_db_index_usage_stats`

A. Option A

B. Option B

C. Option C

D. Option D

Answer: D

NEW QUESTION 126

- (Exam Topic 5)

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

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

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

A. distributed availability groups

B. database mirroring

C. log shipping

D. Database Migration Assistant

Answer: D

Explanation:

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

NEW QUESTION 129

- (Exam Topic 5)

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

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

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

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

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

▼

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

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

▼

a header
a data
an empty

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, table Description automatically generated
 Box 1: CSV files that have file named beginning with "tripdata_2020" Box 2: a header
 FIRSTROW = 'first_row'
 Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.
 Reference:
<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset>

NEW QUESTION 134

- (Exam Topic 5)
 You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.
 Which switch should you use to switch between languages?

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

Answer: B

Explanation:

You can override the default language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.
 Reference:
<https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use>

NEW QUESTION 136

- (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 SQL Server on an Azure virtual machine named SQL1. SQL1 has an agent job to back up all databases.

You add a user named dbadmin1 as a SQL Server Agent operator. You need to ensure that dbadmin1 receives an email alert if a job fails.

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

Actions	Answer Area
Create a job alert	
Create a job notification	
Enable Database Mail	⤵
Enable the email settings for the SQL Server Agent	⤴
Create a job target	⤴ ⤵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Enable the email settings for the SQL Server Agent.

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Step 2: Create a job alert

Step 3: Create a job notification Example:

-- adds an e-mail notification for the specified alert (Test Alert)

-- This example assumes that Test Alert already exists

-- and that François Ajenstat is a valid operator name. USE msdb ;

GO

EXEC dbo.sp_add_notification

@alert_name = N'Test Alert',

@operator_name = N'François Ajenstat',

@notification_method = 1 ; GO

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status> <https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator>

NEW QUESTION 150

- (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 152

- (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 create additional tempdb files. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Reference:

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

NEW QUESTION 154

- (Exam Topic 5)

You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2. You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable query optimizer fixes.
- B. Enable Resource Governor.
- C. Enable parameter sniffing.
- D. Create a workload group.
- E. Configure In-memory OLTP.
- F. Run the Database Engine Tuning Advisor.
- G. Reduce the Max Degree of Parallelism value.

Answer: BC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server>

NEW QUESTION 155

- (Exam Topic 5)

You need to use an Azure Resource Manager (ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O

performance for the SQL Server database and log files

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

```

"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8, "logDisksCount": 1,
    ...
  }
}

"resources": [
  ...
  {
    "osDisk": {
      ...
    },
    "copy": [
      {
        "name": "dataDisks", "count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
        "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",
        "caching": "[if(greaterOrEquals(copyIndex('dataDisks'), parameters('dataDiskCount')),
          variables('dataDisks').caching )]", "diskSizeGB": 1023,
  }
]

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Read only ReadWrite

NEW QUESTION 159

- (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 schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Must use an Azure Data Factory, not an Azure Databricks job. Reference:

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

NEW QUESTION 163

- (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 168

- (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 170

- (Exam Topic 5)

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

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

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


Actions	Answer Area	
Change the Active Directory Admin on TestServer1	➔	⬆
Change the server name and related variables in the templates	➔	⬆
From the database project, deploy the database schema and permissions	➔	⬆
Add IP addresses to the firewall	➔	⬆
From the Azure portal, export the Azure Resource Manager templates	➔	⬆
From the Azure portal, deploy the templates.	➔	⬆

- A. Mastered

B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 172

- (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%',
ENCRYPTION 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 175

- (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 179

- (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 181

- (Exam Topic 5)

You have an Azure SQL database named DB1.

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

Which deployment option and pricing tier should you configure?

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

Answer: A

Explanation:

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.

Reference:

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

NEW QUESTION 184

- (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 189

- (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 194

- (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 198

- (Exam Topic 5)

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

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

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

- A. Yes
- B. No

Answer: A

Explanation:

Reference:

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

NEW QUESTION 200

- (Exam Topic 5)

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

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

Autogrow:

Autoshrink:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 201

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL database. The database fails to respond to queries in a timely manner. You need to identify whether the issue relates to resource_semaphore waits.

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

NOTE: Each correct selection is worth one point.

```
SELECT   

SUM(wait_time) AS total_wait_time_ms  

FROM sys.  

JOIN sys.dm_exec_sessions AS dmv2  

ON dmv1.session_id = dmv2.session_id  

WHERE is_user_process = 1  

GROUP BY wait_type  

ORDER BY SUM(wait_time) DESC;
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Reference:

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

NEW QUESTION 205

- (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 210

- (Exam Topic 5)

You have an Azure SQL database named sqldb1.

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

- A. DBCC SHRINKDATABASE
- B. sp_clean_db_free_space
- C. sp_clean_db_file_free_space
- D. DBCC SHRINKFILE

Answer: A

Explanation:

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

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

NEW QUESTION 214

- (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 216

- (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 218

- (Exam Topic 5)

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

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

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

Answer Area

To prepare for the migration:

To perform the migration, use:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To prepare for the migration:

To perform the migration, use:

NEW QUESTION 222

- (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 225

- (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://rajanireshkaushikk.com/2020/09/09/how-to-choose-right-data-distribution-strategy-for-azure-synapse/>

NEW QUESTION 230

- (Exam Topic 4)

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.

Which Azure Storage functionality should you include in the solution?

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

Answer: C

Explanation:

The lifecycle management policy lets you:

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

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

NEW QUESTION 234

- (Exam Topic 3)

Which windowing function should you use to perform the streaming aggregation of the sales data?

- A. Sliding
- B. Hopping
- C. Session
- D. Tumbling

Answer: D


Explanation:

Scenario: The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping.

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

Timeline Description automatically generated

Tell me the count of Tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun>

NEW QUESTION 239

- (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 240

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 244

- (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 245

- (Exam Topic 1)

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

Collect metrics from:

	▼
The database only	
The elastic pool and the database	
The elastic pool only	
The server, the elastic pool, and the database	

Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: The server, the elastic pool, and the database Scenario: SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Box 2: Azure Event hubs

Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform. Event hubs are able to handle custom metrics.

NEW QUESTION 250

- (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 252

- (Exam Topic 1)

You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure. The solution must meet the business requirements.

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

NOTE: Each correct selection is worth one point.

Quorum model:

	▼
Cloud witness	
Disk witness	
File share witness	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Scenario: Business Requirements

Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.

Box 1: Cloud witness

If you have a Failover Cluster deployment, where all nodes can reach the internet (by extension of Azure), it is recommended that you configure a Cloud Witness as your quorum witness resource.

Box 2: Azure Basic Load Balancer

Microsoft guarantees that a Load Balanced Endpoint using Azure Standard Load Balancer, serving two or more Healthy Virtual Machine Instances, will be available 99.99% of the time.

Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a public Virtual IP (VIP) that is accessible over the internet. An internal listener uses an internal load balancer and only supports clients within the same Virtual Network.

Reference:

<https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness> https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/

NEW QUESTION 257

- (Exam Topic 1)

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

What should you run as part of the implementation?

- A. CREATE LOGIN and the FROM WINDOWS clause
- B. CREATE USER and the FROM CERTIFICATE clause
- C. CREATE USER and the FROM LOGIN clause
- D. CREATE USER and the ASYMMETRIC KEY clause
- E. CREATE USER and the FROM EXTERNAL PROVIDER clause

Answer: E

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.) Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user: CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;

Reference:

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

NEW QUESTION 259

- (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"



* 5. Schedule task (Step 4)

Steps:1. Click on Schedules2. 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 261

- (Exam Topic 1)

You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.

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

Azure Database Migration Service pricing tier:

- Standard 2-vCore
- Standard 4-vCore
- Premium 4-vCore

Required Azure resource:

- A virtual network that has service endpoints
- A VPN gateway
- An Azure Logic app

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure Database Migration service

Box 1: Premium 4-vCore

Scenario: Migrate the SERVER1 databases to the Azure SQL Database platform.

> Minimize downtime during the migration of the SERVER1 databases.

Premium 4-vCore is for large or business critical workloads. It supports online migrations, offline migrations, and faster migration speeds.

Reference: <https://azure.microsoft.com/pricing/details/database-migration/>

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

NEW QUESTION 265

- (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	Answer Area
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	

- 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 269

- (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: ▼

Deployment option: ▼

- 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 274

- (Exam Topic 1)

You create all of the tables and views for ResearchDB1.

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Reference:

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

NEW QUESTION 275

- (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 276

- (Exam Topic 1)

What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?

- A. Configure SalesSQLDb1 as writable.
- B. Update the connection strings of SalesSQLDb1App1.
- C. Update the firewall rules of SalesSQLDb1.
- D. Update the users in SalesSQLDb1.

Answer: C

Explanation:

Scenario: SalesSQLDb1 uses database firewall rules and contained database users.

NEW QUESTION 278

.....

Relate Links

100% Pass Your DP-300 Exam with ExamBible Prep Materials

<https://www.exambible.com/DP-300-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>