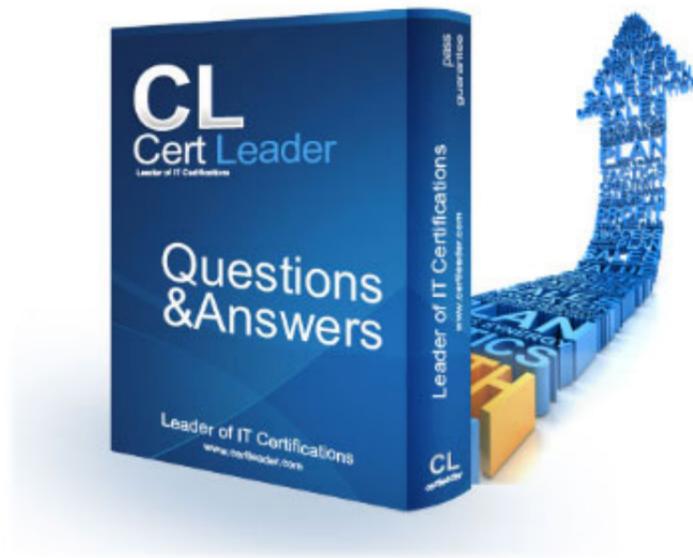


DP-300 Dumps

Administering Relational Databases on Microsoft Azure (beta)

<https://www.certleader.com/DP-300-dumps.html>



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: MAX(<Total number of DBs X average vCore utilization per DB>, <Number of concurrently peaking DBs X Peak vCore utilization per DB>)

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

Reference:

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

NEW QUESTION 2

- (Exam Topic 5)

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

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

NOTE: Each correct selection is worth one point.

BACKUP DATABASE Sales

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

WITH STATS = 5,

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

WITH COPY_ONLY

Box 2: WITH COPY_ONLY

Reference:

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

NEW QUESTION 3

- (Exam Topic 5)

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

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

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

NOTE: Each correct selection is worth one point.

Encryption key: ▼

Encryption type: ▼

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Column encryption Key

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

Reference:

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

NEW QUESTION 4

- (Exam Topic 5)

You 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 5

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table.

You need to identify the extent of the data skew in Table1. What should you do in Synapse Studio?

- A. Connect to Pool1 and query sys.dm_pdw_nodes_db_partition_stats.
- B. Connect to the built-in pool and run DBCC CHECKALLOC.
- C. Connect to Pool1 and run DBCC CHECKALLOC.
- D. Connect to the built-in pool and query sys.dm_pdw_nodes_db_partition_stats.

Answer: D

Explanation:

Use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data. Reference:

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

NEW QUESTION 6

- (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 7

- (Exam Topic 5)

You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1. You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

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

Answer: C

Explanation:

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

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

Reference:

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

NEW QUESTION 8

- (Exam Topic 5)

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

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

NOTE: Each correct selection is worth one point.

CREATE USER [App1] FROM login Windows EXTERNAL PROVIDER

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 9

- (Exam Topic 5)

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

During peak usage, the database will require the following:

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

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

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

Answer: A

Explanation:

Reference:

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

NEW QUESTION 10

- (Exam Topic 5)

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

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

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

- A. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
- B. Run the SET SHOWPLAN_ALL Transact-SQL statement.
- C. Save the actual execution plan.
- D. Enable Query Store for DB1.

Answer: C

Explanation:

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

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

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

NEW QUESTION 10

- (Exam Topic 5)

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

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

NOTE: Each correct selection is worth one point.

- A. Turn on auditing and write audit logs to an Azure Storage account.
- B. Add extended properties to the column.
- C. Turn on Advanced Data Security for the Azure SQL server.
- D. Apply sensitivity labels named Highly Confidential to the column.
- E. Turn on Azure Advanced Threat Protection (ATP).

Answer: ACD

Explanation:

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

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

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

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

Reference:

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

NEW QUESTION 12

- (Exam Topic 5)

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

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

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

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

Answer: D

Explanation:

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

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

Reference:

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

NEW QUESTION 16

- (Exam Topic 5)

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:

The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.

Administrative effort must be minimized. What should you use to host the databases?

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

Answer: D

NEW QUESTION 20

- (Exam Topic 5)

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

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

Which restore option should you use on the secondary replica?

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 24

- (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

Answer Area

```
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];
```



- 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 27

- (Exam Topic 5)

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

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

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

Answer: C

Explanation:

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

What resources the client may access.

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

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

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

Reference:

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

NEW QUESTION 28

- (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 29

- (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 30

- (Exam Topic 5)

You have an on-premises multi-tier application named App1 that includes a web tier, an application tier, and a Microsoft SQL Server tier. All the tiers run on Hyper-V virtual machines.

Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure. You need to recommend a solution to fail over the database tier of App1 to Azure. The solution must provide the ability to test failover to Azure without affecting the current environment.

What should you include in the recommendation?

- A. Azure Backup
- B. Azure Information Protection
- C. Windows Server Failover Cluster
- D. Azure Site Recovery

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure>

NEW QUESTION 31

- (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 36

- (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:

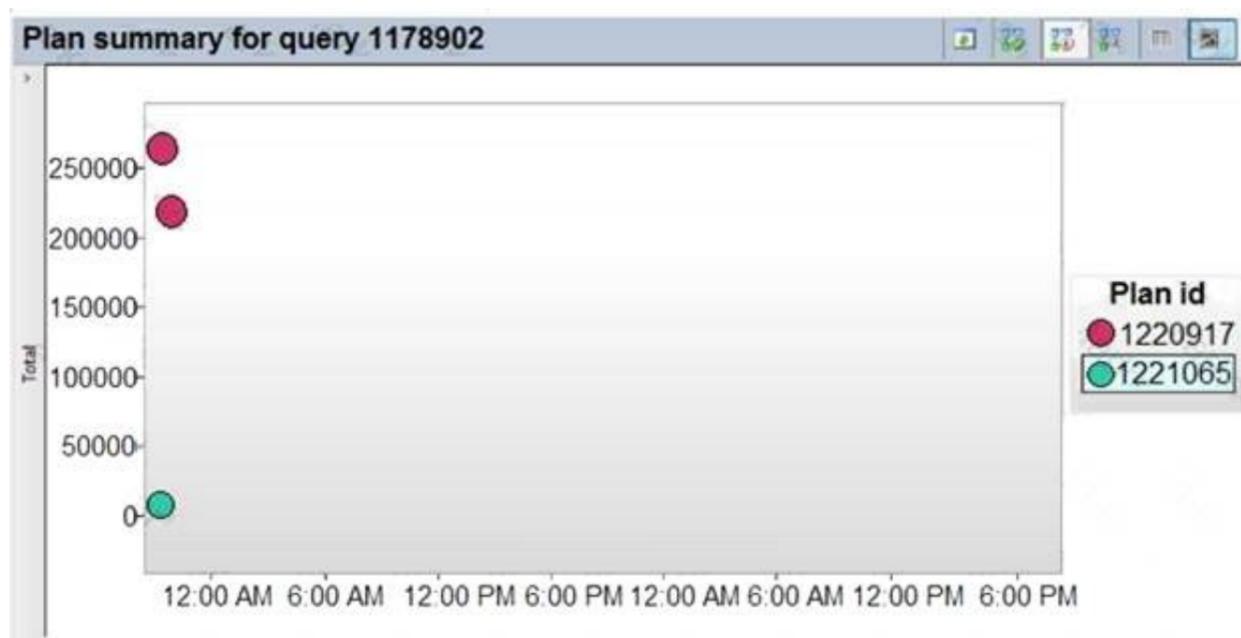
<https://www.sqlshack.com/automate-azure-sql-database-indexes-and-statistics-maintenance/>

NEW QUESTION 37

- (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 40

- (Exam Topic 5)

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

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

You have an Azure Data Lake Storage account that contains a staging zone.

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 45

- (Exam Topic 5)

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

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

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

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

Answer: A

Explanation:

Reference:

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

NEW QUESTION 48

- (Exam Topic 5)

You have the following Azure Resource Manager template.

```

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

```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A screenshot of a computer Description automatically generated with low confidence

Reference:

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

NEW QUESTION 52

- (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 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)

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 59

- (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 61

- (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 62

- (Exam Topic 5)

You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1.

You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server.

Which resource type should you use to create the alert?

- A. Resource Groups
- B. SQL Servers
- C. SQL Databases
- D. SQL Virtual Machines

Answer: C

Explanation:

There are resource types related to application code, compute infrastructure, networking, storage + databases. You can deploy up to 800 instances of a resource type in each resource group.

Some resources can exist outside of a resource group. These resources are deployed to the subscription, management group, or tenant. Only specific resource types are supported at these scopes.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-providers-and-types>

NEW QUESTION 64

- (Exam Topic 5)

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

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

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

Actions

Answer Area

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 66

- (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 67

- (Exam Topic 5)

You have an Azure Data Factory pipeline that is triggered hourly. The pipeline has had 100% success for the past seven days. The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

```
ErrorCode=UserErrorFileNotFound,  
'Type=Microsoft.DataTransfer.Common.Shared.HybridDeliveryException,Message=ADLS  
Gen2 operation failed for: Operation returned an invalid status code  
'NotFound'. Account: 'contosoproduksouth' FileSystem: wwi.Path:  
'BIKES/CARBON/year=2021/month=01/day=10/hour=06'. ErrorCode:  
'PathNotFound'.Message: 'The specified path does not exist.'. RequestId:  
'6d269b78-901f-001b-4924-e7a7bc000000'. Timestamp: 'Sun, 10 Jan 2021 07:45:05
```

What is a possible cause of the error?

- A. From 06:00 to 07:00 on January 10, 2021, there was no data in wwi/BIKES/CARBON.
- B. The parameter used to generate year=2021/month=01/day=10/hour=06 was incorrect.
- C. From 06:00 to 07:00 on January 10, 2021, the file format of data in wwi/BIKES/CARBON was incorrect.
- D. The pipeline was triggered too early.

Answer: B

NEW QUESTION 71

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

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

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

What should you do?

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 76

- (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 78

- (Exam Topic 5)

You have 50 Azure SQL databases.

You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.

What should you do?

- A. Create a diagnostic setting for the activity log that has the Security log enabled.
- B. For the database, create a diagnostic setting that has the InstanceAndAppAdvanced metric enabled.
- C. Create an alert rule that uses a Metric signal type.
- D. Create an alert rule that uses an Activity Log signal type.

Answer: D

Explanation:

Activity log events - An alert can trigger on every event, or, only when a certain number of events occur. Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/alerts-insights-configure-portal>

NEW QUESTION 81

- (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 82

- (Exam Topic 5)

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

- Support failovers that do not require client applications to change their connection strings.
- Replicate the database to a secondary Azure region.
- Support failover to the secondary region. What should you include in the recommendation?

- A. failover groups
- B. transactional replication
- C. Availability Zones
- D. geo-replication

Answer: A

Explanation:

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).

Reference:

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

NEW QUESTION 87

- (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 88

- (Exam Topic 5)

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

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

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

Answer: B

Explanation:

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

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

Reference:

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

NEW QUESTION 89

- (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 91

- (Exam Topic 5)

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

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

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

Answer: BC

Explanation:

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

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

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

NEW QUESTION 96

- (Exam Topic 5)

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

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

What should you include in the recommendation?

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

Answer: B

Explanation:

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

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

Example:

XXXX-XXXX-XXXX-1234

NEW QUESTION 97

- (Exam Topic 5)

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

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

What should you use?

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 100

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 is 30 TB and has a 1-GB daily rate of change.

You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands. You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00.

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

Actions

Answer Area

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

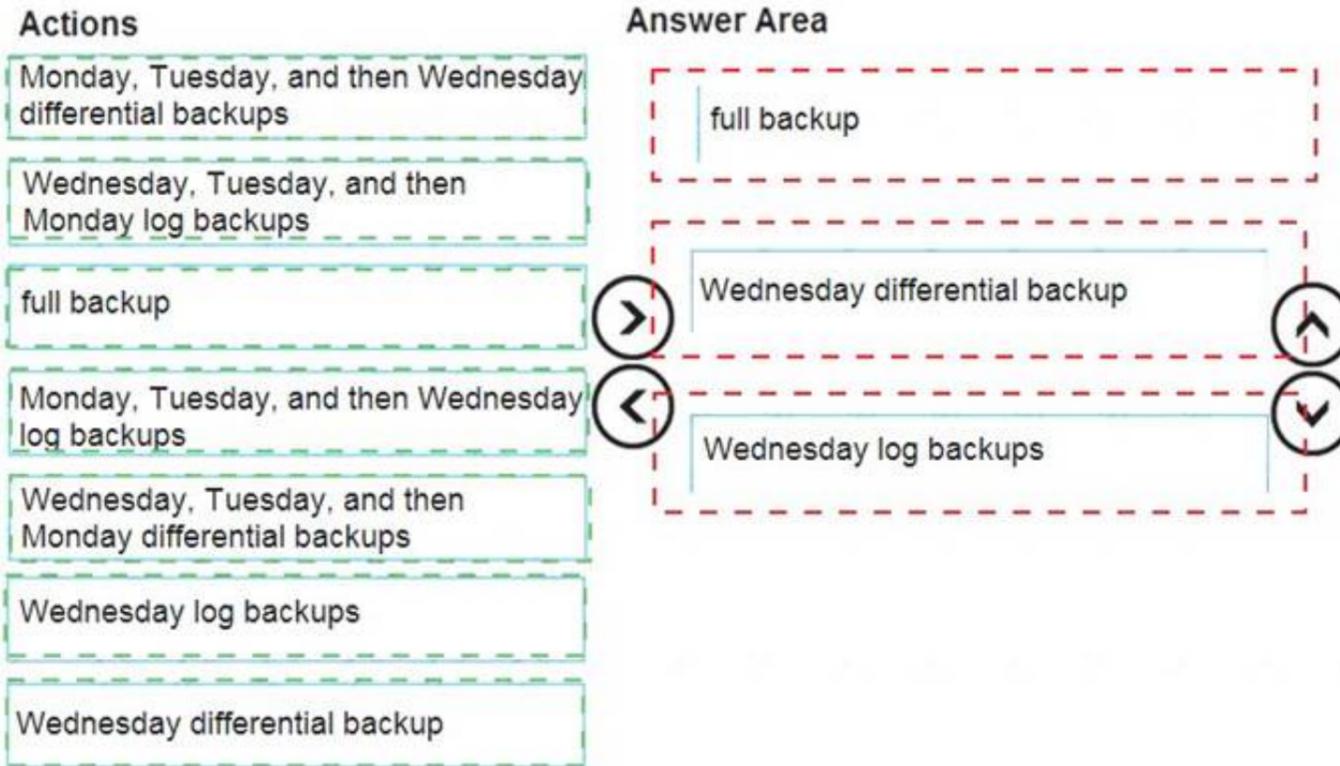
Wednesday differential backup



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 102

- (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:

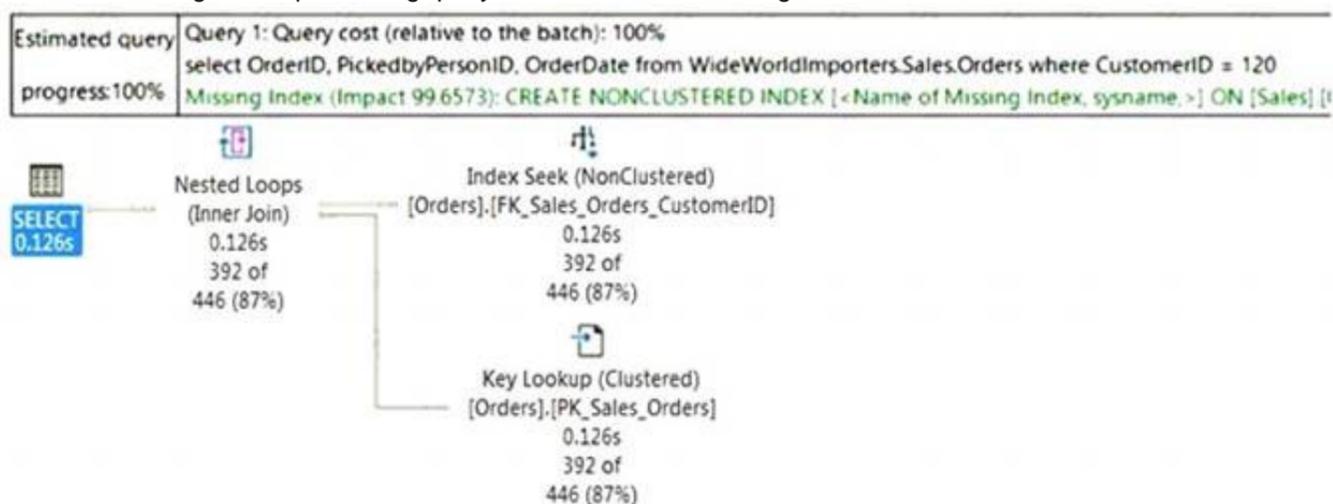
- * 1. Not enough processing resources in Stream Analytics to handle the volume of input events.
- * 2. Not enough throughput within the input event brokers, so they are throttled.
- * 3. Output sinks are not provisioned with enough capacity, so they are throttled. Reference: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling>

NEW QUESTION 105

- (Exam Topic 5)

You have an Azure SQL database.

You are reviewing a slow performing query as shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

The exhibit shows [answer choice].

- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The [answer choice] operator in the execution plan indicates that the query would benefit from performance tuning.

- Index Seek
- Key Lookup
- Nested Loops

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver>

NEW QUESTION 106

- (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 111

- (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 116

- (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 121

- (Exam Topic 5)

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue. You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort. What should you use?

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

Answer: C

Explanation:

Use the Query Store Page in SQL Server Management Studio.

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

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query. This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

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

NEW QUESTION 126

- (Exam Topic 5)

You have an Azure Stream Analytics job.

You need to ensure that the job has enough streaming units provisioned. You configure monitoring of the SU % Utilization metric.

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

- A. Late Input Events
- B. Out of order Events
- C. Backlogged Input Events
- D. Watermark Delay
- E. Function Events

Answer: CD

Explanation:

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

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

Reference:

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

NEW QUESTION 129

- (Exam Topic 5)

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

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

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

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

What should you do?

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

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-managed-instance>**NEW QUESTION 132**

- (Exam Topic 5)

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

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

You have an Azure Data Lake Storage account that contains a staging zone.

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: A**Explanation:**

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

Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>**NEW QUESTION 134**

- (Exam Topic 5)

You are designing a star schema for a dataset that contains records of online orders. Each record includes an order date, an order due date, and an order ship date.

You need to ensure that the design provides the fastest query times of the records when querying for arbitrary date ranges and aggregating by fiscal calendar attributes.

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

- A. Create a date dimension table that has a DateTime key.
- B. Create a date dimension table that has an integer key in the format of YYYYMMDD.
- C. Use built-in SQL functions to extract date attributes.
- D. Use integer columns for the date fields.
- E. Use DateTime columns for the date fields.

Answer: BD**Explanation:**

Reference:

https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-in-a**NEW QUESTION 138**

- (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 142**

- (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 file 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 143

- (Exam Topic 5)
You have a Microsoft SQL Server 2017 server.
You need to migrate the server to Azure. The solution must meet the following requirements:

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

What should you use?

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

Answer: A

NEW QUESTION 146

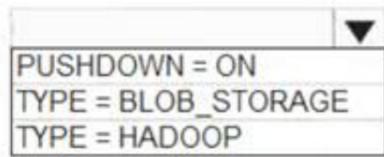
- (Exam Topic 5)
You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.
You plan to access the files in Account1 by using an external table.
You need to create a data source in Pool1 that you can reference when you create the external table. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

CREATE EXTERNAL DATA SOURCE source1

WITH

(LOCATION = 'https://account1.  .core.windows.net',



)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, table Description automatically generated

Box 1: blob

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

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

Box 2: HADOOP

Reference:

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

NEW QUESTION 147

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL database named SQL1. SQL1 is in an Azure region that does not support availability zones. You need to ensure that you have a secondary replica of SQL1 in the same region. What should you use?

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

Answer: C

NEW QUESTION 152

- (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 156

- (Exam Topic 5)

Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1. You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance. Which service tier should you use?

- A. Hyperscale
- B. General Purpose
- C. Premium
- D. Business Critical

Answer: A

NEW QUESTION 158

- (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 161

- (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 165

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

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

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

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

Answer: C

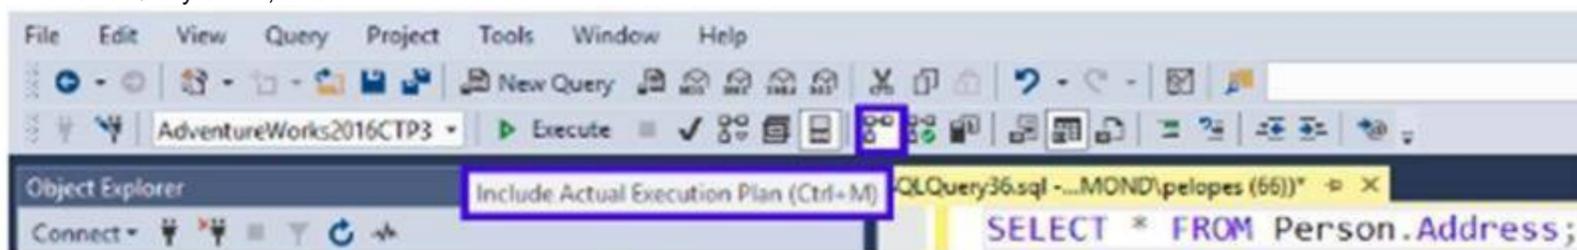
Explanation:

To include an execution plan for a query during execution

* 1. On the SQL Server Management Studio toolbar, click Database Engine Query. You can also open an

existing query and display the estimated execution plan by clicking the Open File toolbar button and locating the existing query.

- * 2. Enter the query for which you would like to display the actual execution plan.
- * 3. On the Query menu, click Include Actual Execution Plan or click the Include Actual Execution Plan toolbar button.



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

Reference:

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

NEW QUESTION 168

- (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 173

- (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 177

- (Exam Topic 5)

You have an Azure subscription.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

```

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

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 178

- (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	<div style="display: flex; align-items: center; justify-content: center;"> <div style="border: 1px solid black; border-radius: 50%; width: 30px; height: 30px; margin-right: 10px;"></div> <div style="font-size: 24px; margin-right: 10px;">></div> <div style="border: 1px solid black; border-radius: 50%; width: 30px; height: 30px; margin-right: 10px;"></div> <div style="font-size: 24px; margin-right: 10px;">^</div> </div>
Add IP addresses to the firewall	<div style="display: flex; align-items: center; justify-content: center;"> <div style="border: 1px solid black; border-radius: 50%; width: 30px; height: 30px; margin-right: 10px;"></div> <div style="font-size: 24px; margin-right: 10px;"><</div> <div style="border: 1px solid black; border-radius: 50%; width: 30px; height: 30px; margin-right: 10px;"></div> <div style="font-size: 24px; margin-right: 10px;">v</div> </div>
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:

Actions	Answer Area
Change the Active Directory Admin on TestServer1	<div style="border: 1px solid blue; padding: 5px; margin-bottom: 5px;">From the Azure portal, export the Azure Resource Manager templates</div> <div style="border: 1px solid blue; padding: 5px; margin-bottom: 5px;">Change the server name and related variables in the templates</div> <div style="border: 1px solid blue; padding: 5px; margin-bottom: 5px;">From the Azure portal, deploy the templates.</div> <div style="border: 1px solid blue; padding: 5px;">From the database project, deploy the database schema and permissions</div>
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.	

NEW QUESTION 182

- (Exam Topic 5)

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

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

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

Actions	Answer Area
Encrypt the Salary column by using the randomized encryption type.	<div style="border: 1px solid gray; height: 100px; width: 100%;"></div>
Create a column encryption key.	
Enable Transparent Data Encryption (TDE).	
Encrypt the Salary column by using the deterministic encryption type.	
Apply a dynamic data mask to the Salary column.	
Create a column master key.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create a column master key

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

Step 2: Create a column encryption key.

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

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

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

Reference:

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

NEW QUESTION 183

- (Exam Topic 5)

You are building an Azure Stream Analytics job to retrieve game data. You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game. How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

SELECT as HighestScore

FROM input TIMESTAMP BY CreatedAt

GROUP BY

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.

Analytic Function Syntax:

TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT

DURATION clause> [<WHEN clause>])

Box 2: Tumbling(minute 5)

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.

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://docs.microsoft.com/en-us/stream-analytics-query/topone-azure-stream-analytics> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun>

NEW QUESTION 186

- (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 190

- (Exam Topic 5)

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

You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create a workload group.	
Create a user-defined classifier function.	⬅️ ⬆️
Modify Resource Governor.	➡️ ⬇️
Create a contained database user.	
Create a resource pool.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text, table Description automatically generated

Reference:

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

NEW QUESTION 195

- (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 198

- (Exam Topic 5)

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

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

What should you recommend?

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

Answer: C

Explanation:

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

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

Reference:

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

NEW QUESTION 199

- (Exam Topic 5)

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

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

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

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

Answer: BD

Explanation:

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

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

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

Reference:

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

NEW QUESTION 201

- (Exam Topic 5)

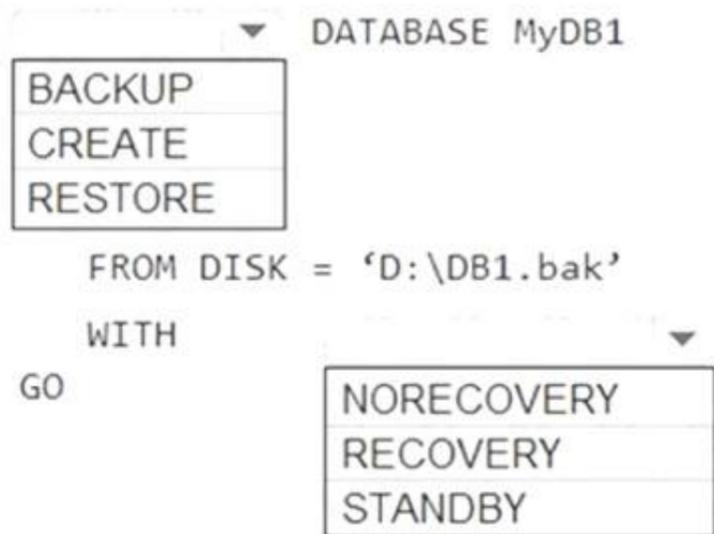
You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named D:\DB1.bak.

You plan to deploy an Always On availability group that will have the following configurations:

- > VM1 will host the primary replica of DB1.
- > VM2 will host a secondary replica of DB1.

You need to prepare the secondary database on VM2 for the availability group.

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondar>

NEW QUESTION 206

- (Exam Topic 5)

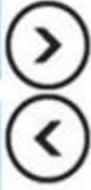
You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.

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

Actions

Answer Area

- Duplicate Azure Policy definitions
- Run Azure Policy remediation tasks
- Create an Azure Blueprints assignment
- Create an Azure Policy initiative
- Create an Azure Policy initiative assignment



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create an Azure Policy Initiative

The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.

With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.

Step 2: Create an Azure Policy Initiative assignment

Assign the initiative definition you created in the previous step. Step 3: Run Azure Policy remediation tasks

To apply the Policy Initiative to the existing SQL databases. Reference:

<https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage>

NEW QUESTION 207

- (Exam Topic 5)

You have an Azure SQL database named db1.

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

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

NOTE: Each correct selection is worth one point.

SELECT *

FROM

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

WHERE database_name = 'db1' AND

start_time >

	▼
DATEADD	
DATEDIFF	
DATEPART	
TODATETIMEOFFSET	

(day, -7, GETDATE())

ORDER BY start_time DESC;

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.resource_stats

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

Box 2: DateAdd

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

DECLARE @s datetime; DECLARE @e datetime;

SET @s= DateAdd(d,-7,GetUTCDate()); SET @e= GETUTCDATE();

SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats

WHERE start_time BETWEEN @s AND @e GROUP BY database_name

HAVING AVG(avg_cpu_percent) >= 80

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-data>

NEW QUESTION 211

- (Exam Topic 5)

You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

 Azure SQL Database built-in intelligence automatically tunes your databases to optimize performance. Click here to learn more about automatic tuning 

Inherit from: 

 The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options 

OPTION	DESIRED STATE	CURRENT STATE
 FORCE PLAN	<input type="button" value="ON"/> <input type="button" value="OFF"/> <input checked="" type="button" value="INHERIT"/>	ON Auto-configured by Azure
 CREATE INDEX	<input type="button" value="ON"/> <input type="button" value="OFF"/> <input checked="" type="button" value="INHERIT"/>	ON Auto-configured by Azure
 DROP INDEX	<input checked="" type="button" value="ON"/> <input type="button" value="OFF"/> <input type="button" value="INHERIT"/>	ON Forced by user

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
Nonclustered indexes will be added to tables to improve performance.	<input type="radio"/>	<input type="radio"/>
Columns will be added to existing indexes automatically.	<input type="radio"/>	<input type="radio"/>
The query execution plan will revert to a previous plan if query performance degrades.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

We see: Tuning option: Create index ON

CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.

Box 2: No

Box 3: Yes

FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.

NEW QUESTION 212

- (Exam Topic 5)

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

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

- A. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (Differential, STATS=5, COMPRESSION);
- B. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (COPY_ONLY, STATS=5, COMPRESSION);
- C. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (File_Snapshot, STATS=5,

COMPRESSION);

D. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (NoInit, STATS=5, COMPRESSION);

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 216

- (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 221

- (Exam Topic 5)

You have an Azure subscription that contains a logical SQL server named Server1. The master database of Server1 contains a user named User1. You need to ensure that User1 can create databases on Server1. Which database role should you assign to User1?

- A. db_owner
- B. dbmanager
- C. dbo
- D. db_ddladmin

Answer: B

NEW QUESTION 222

- (Exam Topic 5)

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

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

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

Answer: DE

Explanation:

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes.

Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference:

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

NEW QUESTION 227

- (Exam Topic 4)

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

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

NOTE: Each correct selection is worth one point.

Table type to store retail store data:

	▼
Hash	
Replicated	
Round-robin	

Table type to store promotional data:

	▼
Hash	
Replicated	
Round-robin	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: Hash Scenario:

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

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

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

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

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

Reference:

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

NEW QUESTION 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 2)

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 237

- (Exam Topic 2)

You are evaluating the role assignments.

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.	<input type="radio"/>	<input type="radio"/>
DBAGroup1 will be able to assign the SQL DB Contributor role to other users.	<input type="radio"/>	<input type="radio"/>
DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

DBAGroup1 is member of the Contributor role.

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

Box 2: No

Box 3: Yes

DBAGroup2 is member of the SQL DB Contributor role.

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

Reference:

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

NEW QUESTION 238

- (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 239

- (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 242

- (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 244

- (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 247

- (Exam Topic 1)

You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals.

What should you include in the recommendation?

- A. For each customer, grant the customer ddl_admin to the existing schema.
- B. For each customer, create an additional schema and grant the customer ddl_admin to the new schema.
- C. For each customer, create an additional schema and grant the customer db_writerto the new schema.
- D. For each customer, grant the customer db_writerto the existing schema.

Answer: D

NEW QUESTION 250

- (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 254

- (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 256

- (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 258

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DP-300 Exam with Our Prep Materials Via below:

<https://www.certleader.com/DP-300-dumps.html>