



Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

NEW QUESTION 1

- (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 2

- (Exam Topic 5)

You have an Azure subscription that uses a domain named contoso.com.

You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.

DBServer1 has a high volume of data changes and low latency requirements for data writes.

You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.

What should you do?

- A. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Asynchronous.
- B. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Synchronous.
- C. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Asynchronous.
- D. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Synchronous.

Answer: C

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on>

NEW QUESTION 3

- (Exam Topic 5)

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

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

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use a serverless SQL pool to create an external table with the extra column.

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

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

NEW QUESTION 4

- (Exam Topic 5)

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

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux.

You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

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

Answer Area

To prepare for the migration:

To perform the migration, use:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To prepare for the migration:

To perform the migration, use:

NEW QUESTION 5

- (Exam Topic 5)

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures.

You need to ensure that SOL Server Agent jobs running on SQL 1 can notify administrators when a failure occurs.

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

- Enable pager notifications upon failure.
- Create a profile named application_dbmail_profile.
- Create a Database Mail account.
- Create a profile named AzureManagedInstance_dbmail_profile.
- Enable email notifications upon failure.

Answer Area

- 1
- 2
- 3

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- Enable pager notifications upon failure.
- Create a profile named application_dbmail_profile.
- Create a Database Mail account.
- Create a profile named AzureManagedInstance_dbmail_profile.
- Enable email notifications upon failure.

Answer Area

- 1 Create a Database Mail account.
- 2
- 3 Create a profile named AzureManagedInstance_dbmail_profile.

Enable email notifications upon failure.

NEW QUESTION 6

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machine named SQL1.
You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort.
Where should you store the metrics?

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

Answer: A

NEW QUESTION 7

- (Exam Topic 5)

You have an Azure subscription.
You need to deploy an Azure SQL database. The solution must meet the following requirements:

- Dynamically scale CPU resources.
- Ensure that the database can be paused to reduce costs. What should you use?

- A. the Business Critical service tier
- B. the serverless compute tier
- C. an elastic pool
- D. the General Purpose service tier

Answer: B

NEW QUESTION 8

- (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: BD

Explanation:

Reference:

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

NEW QUESTION 9

- (Exam Topic 5)

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

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

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

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

Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 10

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:

- Capture raw event data and store the data in Azure Storage.
- Minimize the performance impact of capturing extended events.

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.

Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
ACTION (sqlserver.sql_text)
WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0.
event_file
event_file
event_stream
ring_buffer
)
SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
(MAX_MEMORY = 10 MB,
EVENT_RETENTION_MODE=
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_SINGLE_EVENT_LOSS
NO_EVENT_LOSS
MAX_DISPATCH_LATENCY = 3 SECONDS)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
ACTION (sqlserver.sql_text)
WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0.
event_file
event_file
event_stream
ring_buffer
)
SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
(MAX_MEMORY = 10 MB,
EVENT_RETENTION_MODE=
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_SINGLE_EVENT_LOSS
NO_EVENT_LOSS
MAX_DISPATCH_LATENCY = 3 SECONDS)
```

NEW QUESTION 10

- (Exam Topic 5)

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

- > Send the output to an Azure Synapse.

- > Identify spikes and dips in time series data.
- > Minimize development and configuration effort.

Which should you include in the solution?

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

Answer: C

Explanation:

Anomalies can be identified by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics Reference:
<https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/> <https://docs.microsoft.com/en-us/azure/stream-analytics/azure-synapse-analytics-output>

NEW QUESTION 13

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.

You need to ensure that User1 can perform the following tasks on DB1:

- Create jobs.
- View all jobs.
- Modify, delete, and disable the jobs the user created. The solution must use the principle of least privilege.

Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Location:

- DB1
- master
- msdb

Built-in role:

- SQLAgentOperatorRole
- SQLAgentReaderRole
- SQLAgentUserRole

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<https://learn.microsoft.com/en-us/sql/ssms/agent/sql-server-agent-fixed-database-roles?view=sql-server-ver16#s>

NEW QUESTION 14

- (Exam Topic 5)

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

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

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 18

- (Exam Topic 5)

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

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

You have an Azure SQL database named Sales.

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

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

➤ Ensure that Sales remains available if a datacenter fails.

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

- A. Yes
- B. No

Answer: B

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster. By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

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

NEW QUESTION 21

- (Exam Topic 5)

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

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 23

- (Exam Topic 5)

You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1. You need to identify and repair any physical or logical corruption in DB1. The solution must meet the following requirements:

- Minimize how long it takes to complete the procedure.
- Minimize data loss.

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

Answer Area

DBCC CHECK [DB1],

NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

) WITH

EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

DBCC CHECK [DB1],

NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

) WITH

EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

NEW QUESTION 28

- (Exam Topic 5)

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

What should you do?

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 32

- (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 36

- (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 vCPUs and 64 GB of memory. You need to provide the lowest latency for tempdb.

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

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

Answer: 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 38

- (Exam Topic 5)

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

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

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

Create:

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


Control access to the rows by using:

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 39

- (Exam Topic 5)

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 40

- (Exam Topic 5)

You have an on-premises datacenter that contains a 14-TB Microsoft SQL Server database.

You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance. Which three service tiers support the SQL managed instance? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. General Purpose Standard
- B. Business Critical Premium
- C. Business Critical Memory Optimized Premium
- D. General Purpose Premium
- E. Business Critical Standard

Answer: BCD

NEW QUESTION 43

- (Exam Topic 5)

You are monitoring an Azure Stream Analytics job.

You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero. You need to ensure that the job can handle all the events. What should you do?

- A. Remove any named consumer groups from the connection and use \$default.
- B. Change the compatibility level of the Stream Analytics job.
- C. Create an additional output stream for the existing input stream.
- D. Increase the number of streaming units (SUs).

Answer: D

Explanation:

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 47

- (Exam Topic 5)

You have an Azure Databricks resource.

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

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

Answer: E

Explanation:

Cloud Provider Infrastructure Logs.

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

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

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

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

Reference:

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

NEW QUESTION 51

- (Exam Topic 5)

You are building a database in an Azure Synapse Analytics serverless SQL pool. You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container. Records are structured as shown in the following sample.

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

The records contain two applicants at most.

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

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

NOTE:Each correct selection is worth one point.

▼ applications

CREATE EXTERNAL TABLE
 CREATE TABLE
 CREATE VIEW

```
WITH (
  LOCATION = 'applications/',
  DATA_SOURCE = applications_ds,
  FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
```

▼ (BULK 'https://contoso1.dfs.core.windows.net/applications/year=*/*.parquet',

CROSS APPLY
 OPENJSON
 OPENROWSET

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: CREATE EXTERNAL TABLE

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

Syntax:

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

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

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

Box 2. OPENROWSET
 When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.

Example: AS

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

OPENROWSET(BULK

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

FORMAT='PARQUET') AS [r]


GROUP BY decennialTime, stateName GO

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

NEW QUESTION 56

- (Exam Topic 5)

You have an Azure SQL database named DB1 in the General Purpose service tier. The performance metrics for DB1 are shown in the following exhibit.



You need to reduce the Log IO percentage. The solution must minimize costs. What should you do?

- A. Increase the number of vCores.
- B. Change Recoverymodel to Simple.
- C. Perform a checkpoint operation.
- D. Change Service tier to Business Critical.

Answer: D

NEW QUESTION 57

- (Exam Topic 5)

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

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

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

Answer: C

NEW QUESTION 60

- (Exam Topic 5)

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

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

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 64

- (Exam Topic 5)

You have an Azure SQL database named DB1.

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

During a database audit, you discover that none of the fields contain more than 150 characters. You need to ensure that you can apply page compression to Table1.

What should you do?

- A. Configure the columns as sparse.
- B. Change the column type to varchar(MAX).
- C. Change the column type to varchar(200).
- D. Change the column type to varchar(200).

Answer: D

Explanation:

Reference:

<https://www.sqlshack.com/sql-varchar-data-type-deep-dive/> <https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/>

NEW QUESTION 67

- (Exam Topic 5)

You have a data warehouse in Azure Synapse Analytics.

You need to ensure that the data in the data warehouse is encrypted at rest. What should you enable?

- A. Transparent Data Encryption (TDE)
- B. Advanced Data Security for this database
- C. Always Encrypted for all columns
- D. Secure transfer required

Answer: A

Explanation:

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview>

NEW QUESTION 71

- (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 75

- (Exam Topic 5)

You have an Azure SQL database named DB1 in the General Purpose service tier. You need to monitor DB1 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.

Answer Area

To collect monitoring data, use:

	▼
A virtual machine An Azure function The Azure Monitor agent	

To store monitoring data, create:

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To collect monitoring data, use:

	▼
A virtual machine An Azure function The Azure Monitor agent	

To store monitoring data, create:

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

NEW QUESTION 76

- (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 81

- (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:

Reference:

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

NEW QUESTION 85

- (Exam Topic 5)

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

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

Actions	Answer Area
Create a Database Mail account.	
Enable pager notifications upon failure.	
Create a profile named AzureManagedInstance_dbmail_profile.	
Enable email notifications upon failure.	
Create a profile named application_dbmail_profile.	

> <

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Create a Database Mail account.	Create a Database Mail account.
Enable pager notifications upon failure.	
Create a profile named AzureManagedInstance_dbmail_profile.	Create a profile named AzureManagedInstance_dbmail_profile.
Enable email notifications upon failure.	Enable email notifications upon failure.
Create a profile named application_dbmail_profile.	

> <

NEW QUESTION 86

- (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 90

- (Exam Topic 5)

You have An Azure SQL managed instance.

You need to configure the SQL Server Agent service to email job notifications. Which statement should you execute?

A)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sysadmin_dbmail_profile';
```

B)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'application_dbmail_profile';
```

C)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'AzureManagedInstance_dbmail_profile';
```

D)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sys_dbmail_profile';
```

- A. Option A
- B. Option B

- C. Option C
- D. Option D

Answer: B

NEW QUESTION 94

- (Exam Topic 5)

You have an Azure subscription that contains a storage account named databasebackups. You have an Azure SQL managed instance named DB1. You need to back up DB1 to databasebackups.

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

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE'
SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-02T19:30:08Z&spr=https&sv=2021-06-08&sr=b&sig=B%2FxEYQi0C%4BqyYCeqlHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL = 'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY_ONLY
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE'
SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-02T19:30:08Z&spr=https&sv=2021-06-08&sr=b&sig=B%2FxEYQi0C%4BqyYCeqlHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL = 'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY_ONLY
```

NEW QUESTION 97

- (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 98

- (Exam Topic 5)

You have an Azure subscription that contains 50 instances of SQL Server on Azure Virtual Machines. The instances host 500 Azure SQL databases. You need to ensure that all the databases have the same configuration. The solution must meet the following requirements:

- Auditing must be enabled.
- Azure Defender must be enabled.
- Public network access must be disabled.
- Administrative effort must be minimized.

Which two resources should you create in the subscription? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. an Azure Policy assignment
- B. an Azure Automation account
- C. an Azure Policy initiative
- D. an Azure Automation runbook
- E. an Azure Policy definition

Answer: CE

NEW QUESTION 101

- (Exam Topic 5)

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

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

What should you include in the backup strategy?

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

Answer: B

Explanation:

Azure Backup provides an Enterprise class backup capability for SQL Server on Azure VMs. All backups are stored and managed in a Recovery Services vault. There are several advantages that this solution provides, especially for Enterprises.

Reference:

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

NEW QUESTION 104

- (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 105

- (Exam Topic 5)

You have the following Azure Data Factory pipelines:

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

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

What should you do to schedule the pipelines for execution?

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 110

- (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	<div style="display: flex; align-items: center; justify-content: center;"> <div style="margin-right: 10px;">></div> <div style="border: 1px solid #ccc; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center;">^</div> </div>
Turn on the TRUSTWORTHY property and the CLR property	<div style="display: flex; align-items: center; justify-content: center;"> <div style="margin-right: 10px;"><</div> <div style="border: 1px solid #ccc; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center;">v</div> </div>
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 112

- (Exam Topic 5)

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

You scale up the virtual machine to 8 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: C

Explanation:

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

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

NEW QUESTION 113

- (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 115

- (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 118

- (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.

FROM DISK = 'D:\DB1.bak'

 WITH

 GO

- 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 123

- (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 128

- (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 129

- (Exam Topic 5)

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1.

You need to ensure that app1 can access db1. The solution must use the principle of least privilege. 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.

Answer Area

CREATE

CREDENTIAL
 LOGIN
 USER

[app1]


FOR LOGIN app1
 FROM EXTERNAL PROVIDER
 FROM LOGIN app1
 WITHOUT LOGIN

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 132

- (Exam Topic 5)

You have an Azure subscription that contain an Azure SQL managed instance named SQLMI1 and a Log Analytics workspace named Workspace1. You need to collect performance metrics for SQLMI1 and stream the metrics to Workspace1.

- A. Create the private endpoint connection on SQLMI1.
- B. Configure Azure SQL Analytics to use Workspace1.
- C. Modify the Computer + storage settings for SQLMI1.
- D. Modify the diagnostic settings for SQLMI1.

Answer: B

NEW QUESTION 133

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1. You need to move a SQL Server Agent job from DB1 to SQLMI1. Which job attribute is unsupported in SQLMI1?

- A. log to table
- B. email notifications
- C. schedules
- D. output files

Answer: D

NEW QUESTION 135

- (Exam Topic 5)

You have an Azure subscription.

You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.

You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs. The solution must minimize administrative effort.

Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

MSX: ▼

SQL database
 SQL managed instances
 SQL virtual machines

TSX: ▼

SQL database
 SQL managed instances
 SQL virtual machines

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

MSX: ▼

- SQL database
- SQL managed instances
- SQL virtual machines

TSX: ▼

- SQL database
- SQL managed instances
- SQL virtual machines

NEW QUESTION 138

- (Exam Topic 5)

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

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

FactSales has 6 billion rows and is loaded nightly by using a batch process. Which type of compression provides the greatest space reduction for the database?

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

Answer: D

Explanation:

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

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

Reference:

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

NEW QUESTION 140

- (Exam Topic 5)

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

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

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

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

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

Answer: C

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

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

NEW QUESTION 141

- (Exam Topic 5)

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

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

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

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

Solution: You run theRemove-AzSqlDatabasePowerShell cmdlet for Database1 on Server2. You run theRestore-AzSqlDatabasePowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

Reference:

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

NEW QUESTION 145

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication. You need to migrate SQL1 to Azure. Which service should you use?

- A. Azure SQL Edge
- B. Azure SQL Database
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed instance

Answer: C

NEW QUESTION 148

- (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 152

- (Exam Topic 5)

You have a SQL Server on Azure Virtual Machines instance named VM1 . You run the following query.

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 155

- (Exam Topic 5)

You configure backups for an Azure SQL database as shown in the following exhibit.

Point-in-time-restore

Specify how long you want to keep your point-in-time backups. [Learn more](#)

How many days would you like PITR backups to be kept?

Long-term retention

Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more](#)

Weekly LTR Backups

Keep weekly backups for:

Monthly LTR Backups

Keep the first backup of each month for:

Yearly LTR Backups

Keep an annual backup for:

Which weekly backup of the year would you like to keep?

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.

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

After the 52nd weekly backup runs, there will be [answer choice]

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

NEW QUESTION 160

- (Exam Topic 5)

You have a database named db1.

The log for db1 contains the following entry.

```
Date 10/5/2021 10:57:08 AM
Log SQL Server (Current - 10/5/2021 11:26:00 AM)

Source spid1595

Message
The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'
```

You need to ensure That db1 can process transactions.

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

Answer Area

- Remove db1 from the availability group.
- Shrink the transaction log file.
- Add db1 back to the availability group.

NEW QUESTION 165

- (Exam Topic 5)

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

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

NOTE: Each correct selection is worth one point.

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

Answer: BE

Explanation:

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

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

NEW QUESTION 168

- (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 an Azure Databricks notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 169

- (Exam Topic 5)

You manage an enterprise data warehouse in Azure Synapse Analytics.

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

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

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

Answer: A

Explanation:

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

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

Reference:

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

NEW QUESTION 174

- (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 177

- (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 182

- (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 184

- (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 187

- (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 189

- (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 193

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named Appl. Appl uses db1. You need to enable Resource Governor for a App1. The solution must meet the following requirements: App1 must be able to consume all available CPU resources.

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

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

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

Actions

- Create a plan.
- Create a classifier function in db1.
- Create a workload group.
- Create a classifier function in the master database.
- Create a resource pool that has the following configurations.

```

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50
            
```

➤
➤

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- Create a plan.
- Create a classifier function in db1.
- Create a workload group.
- Create a classifier function in the master database.
- Create a resource pool that has the following configurations.

```

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50
            
```

➤
➤

Answer Area

- Create a resource pool that has the following configurations.

```

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50
            
```

- Create a workload group.
- Create a classifier function in the master database.

NEW QUESTION 196

- (Exam Topic 5)

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

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

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

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

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

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

NEW QUESTION 201

- (Exam Topic 5)

You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1. You need to configure SQL1 to use mixed mode authentication. Which procedure should you run?

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

Answer: B

NEW QUESTION 205

- (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 209

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil.


You plan to perform an online migration of db1 to MM by using Azure Database Migration Service.

You need to create the backups for the migration. The solution must minimize the number of backup files created.

Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 214

- (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: A

Explanation:

<https://www.mssqltips.com/sqlservertip/3685/live-query-statistics-in-sql-server-2016/>

NEW QUESTION 217

- (Exam Topic 5) You have SQL Server on an Azure virtual machine. You review the query plan shown in the following exhibit.

Shows top 25 resource consumers for database AdventureWorks2017.
 Total logical reads: 5 = 200000, 4 = 200000, 3 = 60000, 2 = 60000
 Key Lookup (Clustered)
 Physical Operation = Key Lookup
 Logical Operation = Key Lookup
 Estimated Execution Mode = Row
 Storage = RowStore
 Estimated Operator Cost = 0.149639 (98%)
 Estimated I/O Cost = 0.003125
 Estimated Subtree Cost = 0.149639
 Estimated CPU Cost = 0.0001581
 Estimated Number of Executions = 48
 Estimated Number of Rows = 1
 Estimated Row Size = 31 B
 Ordered = True
 Node

Key Lookup (Clustered)
 Uses a supplied clustering key to lookup on a table that has a clustered index.

Physical Operation	Key Lookup
Logical Operation	Key Lookup
Estimated Execution Mode	Row
Storage	RowStore
Estimated Operator Cost	0.149639 (98%)
Estimated I/O Cost	0.003125
Estimated Subtree Cost	0.149639
Estimated CPU Cost	0.0001581
Estimated Number of Executions	48
Estimated Number of Rows	1
Estimated Row Size	31 B
Ordered	True
Node ID	8

Object
 [AdventureWorks2017].[Sales].[SalesOrderHeader].
 [PK_SalesOrderHeader_SalesOrderID]

Output List
 [AdventureWorks2017].[Sales].
 [SalesOrderHeader].SubTotal, [AdventureWorks2017].
 [Sales].[SalesOrderHeader].TaxAmt,
 [AdventureWorks2017].[Sales].
 [SalesOrderHeader].Freight

Seek Predicates
 Seek Key(1): Prefix: [AdventureWorks2017].[Sales].
 [SalesOrderHeader].SalesOrderID = Scalar Operator
 ([AdventureWorks2017].[Sales].[SalesOrderHeader].
 [SalesOrderID])

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

Statements

Yes

No

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Reference:

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

NEW QUESTION 220

- (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 224

- (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.

<input type="text"/>
<input type="text" value="CSV files in the tripdata_2020 subfolder"/>
<input type="text" value="files that have files names beginning with 'tripdata_2020'"/>
<input type="text" value="CSV files that have file names containing 'tripdata_202'"/>
<input type="text" value="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.

<input type="text"/>
<input type="text" value="a header"/>
<input type="text" value="a data"/>
<input type="text" value="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 226

- (Exam Topic 5)

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

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

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

Actions	Answer Area
Encrypt the Salary column by using the randomized encryption type.	
Create a column encryption key.	
Enable Transparent Data Encryption (TDE).	<input type="button" value="➤"/>
Encrypt the Salary column by using the deterministic encryption type.	<input type="button" value="⬅"/>
Apply a dynamic data mask to the Salary column.	<input type="button" value="⬆"/>
Create a column master key.	<input type="button" value="⬇"/>

- 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 229

- (Exam Topic 5)

You configure backup for an Azure SQL database as shown in the following exhibit.

Point-in-time-restore
 Specify how long you want to keep your point-in-time backups. [Learn more](#)

How many days would you like PITR backups to be kept?

Long-term retention
 Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more](#)

Weekly LTR Backups
 Keep weekly backups for:

Monthly LTR Backups
 Keep the first backup of each month for:

Yearly LTR Backups
 Keep an annual backup for:

Which weekly backup of the year would you like to keep?

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

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a **[answer choice]**

- point-time restore (PITR) backup
- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.

- 65 backup copies
- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a **[answer choice]**

- point-time restore (PITR) backup.
- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.

- 65 backup copies
- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

NEW QUESTION 234

- (Exam Topic 5)

You have a database on a SQL Server on Azure Virtual Machines instance. The current state of Query Store for the database is shown in the following exhibit.

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text Description automatically generated

NEW QUESTION 239

- (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": [
        "Microsoft.Sql/servers",
        "Microsoft.SqlVirtualMachines/sqlVirtualMachines",
        "Microsoft.Synapse/workspaces/sqldatabases"
      ],
      "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 243

- (Exam Topic 5)

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

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

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

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

NOTE:Each correct selection is worth one point.

Values

Answer Area

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: HASH

Box 2: OrderDateKey

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

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

Reference:

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

NEW QUESTION 247

- (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 250

- (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 252

- (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 257

- (Exam Topic 5)

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

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

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

What should you do first?

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

Answer: D

Explanation:

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

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

Server VM resource provider. Reference:

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

NEW QUESTION 261

- (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 262

- (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: D

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 265

- (Exam Topic 5)

You have an Azure SQL database named DB3.

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



Which Transact-SQL command should you run?

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

Answer: C

Explanation:

The exhibit displays Database [State] properties.

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

Reference:

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

NEW QUESTION 268

- (Exam Topic 5)

You have an Azure subscription that contains the following resources:

- 10 Azure SQL databases
- Five Azure SQL managed instances
- Five instances of SQL Server on Azure Virtual Machines

You need to implement a centralized monitoring solution for all the Azure SQL resources. The solution must minimize administrative effort. What should you include in the solution?

- A. Log Analytics
- B. Azure SQL Analytics
- C. Query Performance Insight
- D. SQL Insights

Answer: B

NEW QUESTION 272

- (Exam Topic 5)

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

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

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance. Solution: You shrink the transaction log file. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

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

NEW QUESTION 275

- (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 Business Critical service tier and Availability Zones.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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 276

- (Exam Topic 5)

You have an Azure SQL database named DB1. A user named User 1 has an Azure AD account.

You need to provide User1 with the ability to add and remove columns from the tables inDBV The solution must use the principle of least privilege.

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

- A. Assign the database user the db.ddladmm role.
- B. Assign the database user the db.owner role.
- C. Create a contained database user.
- D. Create a login and an associated database user.

Answer: AD

NEW QUESTION 277

- (Exam Topic 5)

You have an Azure subscription.

You plan to deploy an instance of SQL Server on Azure Virtual Machines that supports Write Accelerator.

Which virtual machine series should you use?

- A. H-series
- B. G-series
- C. M-series
- D. E-series

Answer: C

NEW QUESTION 281

- (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 285

- (Exam Topic 5)

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

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

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

- Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- Restrict network connectivity to SqlSrv1. What should you create on VNet1?

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

Answer: B

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

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

NEW QUESTION 288

- (Exam Topic 5)

You have an Azure SQL managed instance.

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

What should you do?

- A. Generate an estimated execution plan.
- B. Generate an actual execution plan.
- C. Run sys.dm_exec_query_plan_stats.
- D. Generate Live Query Statistics.

Answer: C

Explanation:

Reference:

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

NEW QUESTION 291

- (Exam Topic 5)

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

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

```

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

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

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Read only ReadWrite

NEW QUESTION 295

- (Exam Topic 5)

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

- Support real-time data replication to a different geographic region.
- Use Azure as a disaster recovery target.
- Minimize costs and administrative effort.

What should you include in the recommendation?

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

Answer: D

NEW QUESTION 298

- (Exam Topic 5)

You have an on-premises datacenter that contains a 2-TB Microsoft SQL Server 2019 database named DB1. You need to recommend a solution to migrate DB1 to an Azure SQL managed instance. The solution must minimize downtime and administrative effort.

What should you include in the recommendation?

- A. Log Replay Service (LRS)
- B. log shipping
- C. transactional replication
- D. SQL Data Sync

Answer: B

NEW QUESTION 302

- (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. database mirroring
- B. distributed availability groups
- C. Always On Availability Group
- D. Azure Database Migration Service

Answer: D

NEW QUESTION 305

- (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 307

- (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 311

- (Exam Topic 5)

You configure a long-term retention policy for an Azure SQL database as shown in the exhibit. (Click the Exhibit tab.)

Configure policies ✕

SQL server

Point in Time Restore Configuration

Configure PiTR backup retention ▼ Days

Long-term Retention Configurations

Weekly LTR Backups ⓘ

How long would you like weekly backups to be kept?

6 Week(s) ▼

Monthly LTR Backups ⓘ

How long would you like the first backup of each month to be kept?

12 Month(s) ▼

Yearly LTR Backups ⓘ

Which weekly backup of the year would you like to retain?

Week 2 ▼

How long would you like this annual backup to be kept?

10 Year(s) ▼

The first weekly backup occurred on January 4, 2020. The dates for the first 10 weekly backups are:

- January 4, 2020
- January 11, 2020
- January 18, 2020
- January 25, 2020
- February 1, 2020
- February 8, 2020
- February 15, 2020
- February 22, 2020
- February 29, 2020
- March 7, 2020

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 backup saved to long-term retention on January 4, 2020, will be retained for

▼

6 weeks

12 months

10 years

The backup saved to long-term retention on January 11, 2020 will be retained for

▼

6 weeks

12 months

10 years

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 314

- (Exam Topic 5)

You have an Azure SQL database named DBI that contains a nonclustered index named index1. End users report slow queries when they use index1. You need to identify the operations that are being performed on the index. Which dynamic management view should you use?

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

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

Answer: D

NEW QUESTION 319

- (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. PARAMETERIZATIONtoFORCED
- B. PARAMETERIZATIONtoSIMPLE
- C. Delayed Durability toForced
- D. READ_COMMITTED_SNAPSHOTtoON

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 321

- (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 324

- (Exam Topic 5)

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

The data contains the following columns:

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

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

EventCategory:

▼
DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping:

▼
DimChannel
DimDate
DimEvent
FactEvents

TotalEvents:

▼
DimChannel
DimDate
DimEvent
FactEvents

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, application, table Description automatically generated

Box 1: FactEvents

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

Box 2: DimChannel

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

Box 3: DimEvent Reference:

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

NEW QUESTION 329

- (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 334

- (Exam Topic 4)

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

What should you create?

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

Answer: B

Explanation:

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

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

Reference:

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

NEW QUESTION 337

- (Exam Topic 3)

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

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

Answer: D


Explanation:

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

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

Timeline Description automatically generated

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



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

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

NEW QUESTION 341

- (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 342

- (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 343

- (Exam Topic 2)
 What should you implement to meet the disaster recovery requirements for the PaaS solution?

- A. Availability Zones
- B. failover groups
- C. Always On availability groups
- D. geo-replication

Answer: B

Explanation:

Scenario: In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.

The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active geo-replication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.

The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.

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

NEW QUESTION 348

- (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:
 DTU
 vCore

Deployment option:
 An Azure SQL Database managed instance
 A SQL Server Always On availability group

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: DTU

Scenario:

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

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.

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

Box 2: An Azure SQL database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

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

NEW QUESTION 352

- (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: B

Explanation:

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

NEW QUESTION 357

- (Exam Topic 1)

You create all of the tables and views for ResearchDB1.

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Reference:

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

NEW QUESTION 358

- (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 362

- (Exam Topic 1)

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

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

Actions

Answer Area

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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


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

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

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

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

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



- * 5. Schedule task (Step 4)

Steps:1. Click on Schedules2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.

Reference:

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

NEW QUESTION 363

- (Exam Topic 1)

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

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

NOTE:Each correct selection is worth one point.

- A. sys.dm_pdw_nodes_tran_locks
- B. sys.dm_exec_compute_node_errors
- C. sys.dm_exec_requests
- D. sys.dm_cdc_errors
- E. sys.dm_pdw_nodes_os_wait_stats
- F. sys.dm_tran_locks

Answer: AE

Explanation:

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

A: Use sys.dm_pdw_nodes_tran_locks instead of sys.dm_tran_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information. SELECT

t1.resource_type, t1.resource_database_id, t1.resource_associated_entity_id, t1.request_mode, t1.request_session_id, t2.blocking_session_id

FROM sys.dm_tran_locks as t1

INNER JOIN sys.dm_os_waiting_tasks as t2

ON t1.lock_owner_address = t2.resource_address;

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

sys.dm_os_wait_stats: for SQL Server sys.dm_db_wait_stats: for Azure SQL Database

sys.dm_pdw_nodes_os_wait_stats: for Azure SQL Data Warehouse Reference:

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

NEW QUESTION 366

- (Exam Topic 1)

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

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

NOTE:Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Scenario: Business Requirements

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

Box 1: Cloud witness

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

Box 2: Azure Basic Load Balancer

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

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

Reference:

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

NEW QUESTION 368

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

DP-300 Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The DP-300 Practice Test Here](#)