Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

About Exambible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

   All examinations will be up to date.

* 24/7 Quality Support

   We will provide service round the clock.

* 100% Pass Rate

   Our guarantee that you will pass the exam.

* Unique Gurantee

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

NEW QUESTION 1
- (Exam Topic 5)
You have 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:

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)
You have an Azure subscription.
You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance.
You need to configure the disks on the virtual machine. The solution must meet the following requirements:
• Minimize latency for transaction logs.
• Minimize the impact on IO Of the virtual machine.
Which type of disk should you use for each workload? To answer. drag the appropriate disk types to the correct workloads. Each disk type may be used once, 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.

Disk Types

Local

Premium SSD

Standard HDD

Standard SSD

Ultra Disk

Answer Area

TempDB: [Disk Type]

Transaction logs: [Disk Type]

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated

NEW QUESTION 4
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.
You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.
You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.
What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Run: [▼]

A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in: [▼]

A Recovery Services vault
An Azure Blob storage account
An SMB file share

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, application, Word Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-managed-instance-online

NEW QUESTION 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:

NEW QUESTION 6
- (Exam Topic 5)
Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.
You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions in the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 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 are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.
You need to create a surrogate key for the table. The solution must provide the fastest query performance. What should you use for the surrogate key?

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

Answer: A

Explanation:
Dedicated SQL pool supports many, but not all, of the table features offered by other databases. Surrogate keys are not supported. Implement it with an Identity column.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesovervie

NEW QUESTION 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 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

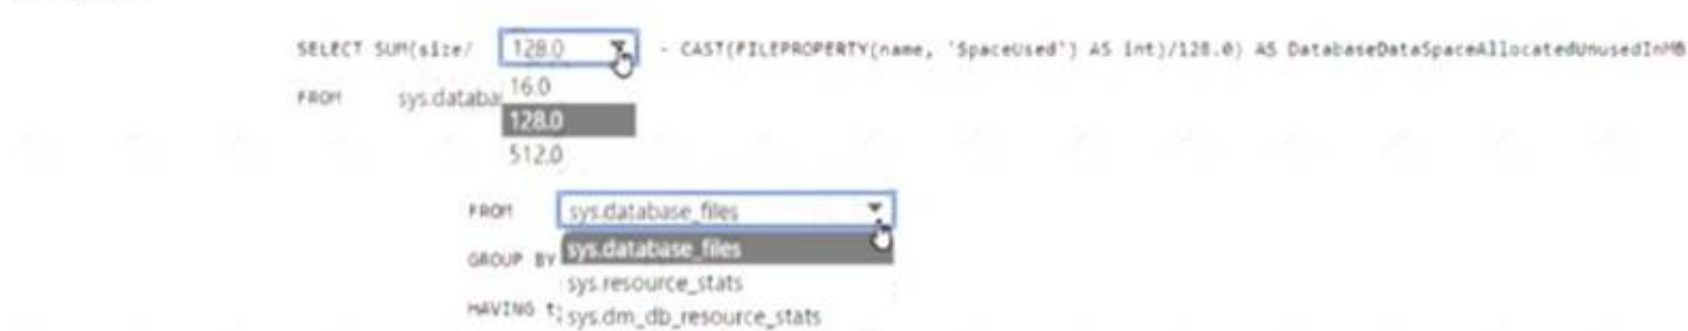**NEW QUESTION 10**
- (Exam Topic 5)
You have an Azure SQL database named D61.
You need to identify how much unused space in megabytes was allocated to DB1.
How should you complete the Transact-SQL query? To answer select the appropriate options m the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

A. Mastered
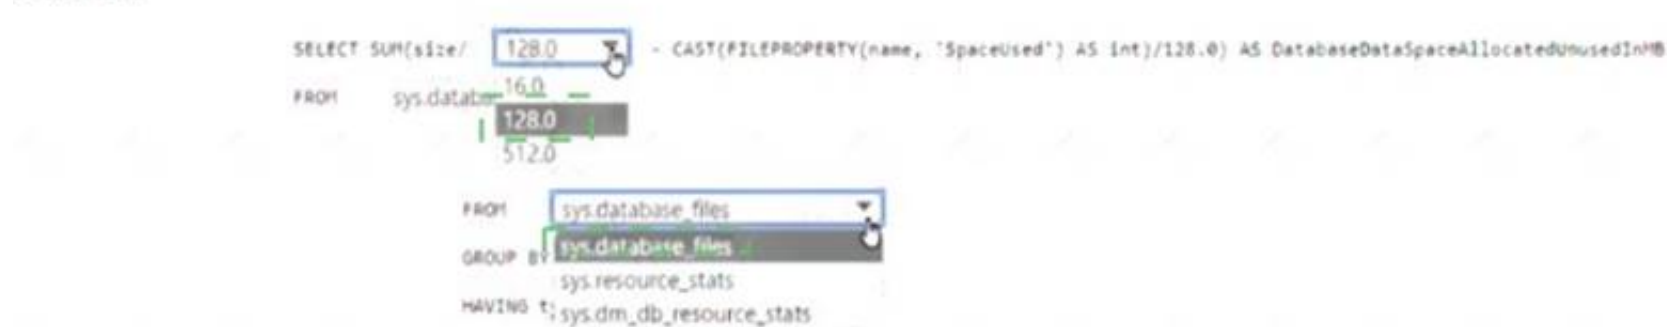B. Not Mastered

Answer: A

Explanation:
Answer Area

NEW QUESTION 10
- (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 14
- (Exam Topic 5)
You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.
You scale up the virtual machine to 16 vCPUSs and 64 GB of memory. You need to provide the lowest latency for tempdb.
What is the total number of data files that tempdb should contain?

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

Answer: D

Explanation:
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical
processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database

NEW QUESTION 16
- (Exam Topic 5)
You 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 18
- (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 22

- (Exam Topic 5)
You have an Azure SQL Database managed instance. The instance starts experiencing performance issues.
You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.
What should you use?

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

Answer: D

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

NEW QUESTION 27
- (Exam Topic 5)
You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.
Which resource provider should you enable?

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

Answer: B

Explanation:
Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger

NEW QUESTION 28
- (Exam Topic 5)
You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in fiveminute 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 33
- (Exam Topic 5)
You have the following Transact-SQL query.

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

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

A. ColumnA
B. ColumnB

C. ColumnC
D. ColumnD

Answer: C

Explanation:

Example:
Free space for the file in the below query result set will be returned by the FreeSpaceMB column. SELECT DB_NAME() AS DbName,
name AS FileName,
type_desc,
size/128.0 AS CurrentSizeMB,
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB FROM sys.database_files
WHERE type IN (0,1);
Reference:
https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/

NEW QUESTION 36
- (Exam Topic 5)
You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

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

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

≫ The first six numerals of each customer's phone number must be masked.

≫ The last four digits of each customer's phone number must be visible.

≫ Hyphens must be preserved and displayed.

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

Exposed Prefix:
```
0
1
3
5
```

Padding String:
```
x
xxxxxx
xxx-xxx
xxx-xxx-
x[3]-x[3]
```

Exposed Suffix:
```
0
1
3
5
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: 0
Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix
Box 2: xxx-xxx
Box 3: 5 Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking

NEW QUESTION 39
- (Exam Topic 5)
You have an Azure SQL database named DB1 in the General Purpose service tier. You need to monitor D81 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 41
- (Exam Topic 5)
You have an Azure SQL database.
Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.
You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.
What should you use?

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

Answer: C

Explanation:
Use the Query Store Page in SQL Server Management Studio.
Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.
Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query.
This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu

NEW QUESTION 45
- (Exam Topic 5)
You have a single availability set that contains two SQL Server on Azure Virtual Machines instances.
The instances were deployed by using an Azure Marketplace SQL Server 2019 Enterprise image that has the latest cumulative updates applied. The instances are configured as the nodes of a failover cluster instance (FCI) named FCI1.
You need to ensure that client applications can connect to FCI1. The solution must meet the following requirements:
• Provide an availability SLA
• Minimize costs.
What should you create?

A. a virtual network name (VNN) resource
B. a Basic Azure Load Balancer
C. a distributed network name (DNN) resource
D. an Azure Standard Load Balancer

Answer: C

NEW QUESTION 50
- (Exam Topic 5)
You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. DF1 contains a linked service.
You have an Azure Key vault named vault1 that contains an encryption kay named key1. You need to encrypt df1 by using key1.
What should you do first?

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

Answer: B

Explanation:
A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.
Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key

NEW QUESTION 51
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

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

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

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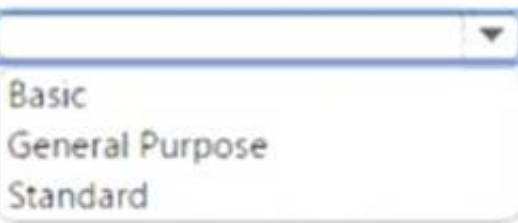
Answer: D

NEW QUESTION 56
- (Exam Topic 5)
You have a Microsoft SQL Server 2017 server that hosts five databases. You Plan to migrate the databases to Azure.
You need to recommend a solution that meets the following requirements:
≫ Automatically scales compute based on the workload demand
≫ Provides per-second billing
What should you include in the recommendation? To answer, select the appropriate options in the answer area.

Azure service:

A single Azure SQL database in the provisioned compute tier
A single Azure SQL database in the serverless compute tier
An Azure SQL Database elastic pool
Azure SQL Managed Instance

Service tier:

Basic
General Purpose
Standard

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Azure service:

A single Azure SQL database in the provisioned compute tier
A single Azure SQL database in the serverless compute tier
An Azure SQL Database elastic pool
Azure SQL Managed Instance

Service tier:

Basic
General Purpose
Standard

NEW QUESTION 57
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1. You need to modify the MAXDOP settings for db1.
What should you do?

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

Answer: D

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

NEW QUESTION 60
- (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 65
- (Exam Topic 5)
You have an Azure SQL Database instance named DatabaseA on a server named Server1.
You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.
You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

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

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

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

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

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

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'
Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:
CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command.
Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]
Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:
CREATE USER readonlyuser FROM LOGIN readonlylogin;
Step 3: On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]
Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to the database via the db_datareader role.
EXEC sp_addrolemember 'db_datareader', 'readonlyuser'; Reference:
https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/

NEW QUESTION 68
- (Exam Topic 5)
You have an Azure Synapse Analytics dedicated SQL pool.
You runPDW_SHOWSPACEUSED('dbo.FactInternetSales');and get the results shown in the following table.

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

Which statement accurately describes thedbo.FactInternetSalestable?

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

Answer: D

Explanation:
The rows per distribution can vary up to 10% without a noticeable impact on performance. Here the distribution varies more than 10%. It is skewed.
Note: SHOWSPACEUSED displays the number of rows, disk space reserved, and disk space used for a specific table, or for all tables in a Azure Synapse Analytics or Parallel Data Warehouse database.
This is a very quick and simple way to see the number of table rows that are stored in each of the 60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.
ROUND_ROBIN distributed tables should not be skewed. Data is distributed evenly across the nodes by design.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distrib
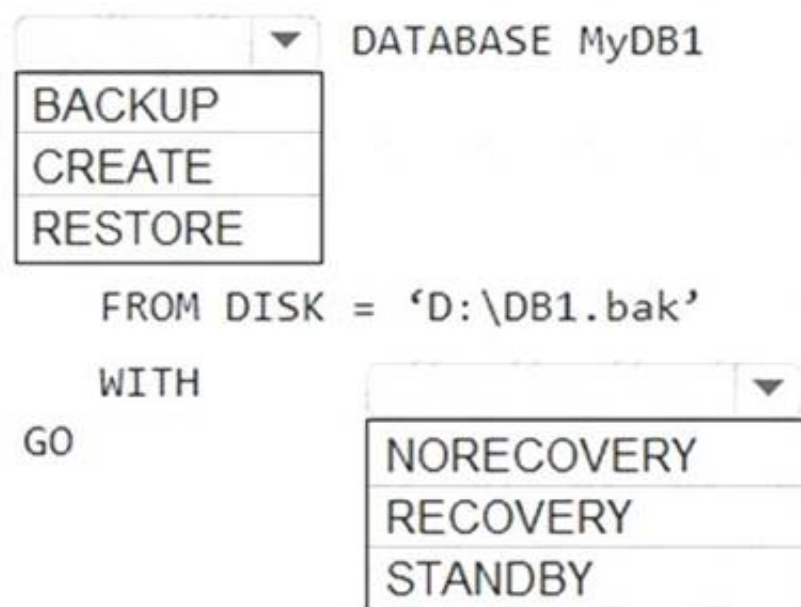
NEW QUESTION 72
- (Exam Topic 5)
You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named D:\DB1.bak.
You plan to deploy an Always On availability 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.

```
[  ▼ ]  DATABASE MyDB1
 BACKUP
 CREATE
 RESTORE

    FROM DISK = 'D:\DB1.bak'

    WITH          [         ▼ ]
GO            NORECOVERY
              RECOVERY
              STANDBY
```

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 77
- (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 sen/ice principal named appl.
You need to ensure that appl 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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

NEW QUESTION 82
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. You have an application that queries DB1 to generate a sales report.
You need to see the parameter values from the last time the query was executed.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

A. EnableLast_Query_Plan_Statsin the master database
B. EnableLightweight_Query_Profilingin DB1
C. EnableLast_Query_Plan_Statsin DB1
D. EnableLightweight_Query_Profilingin the master database
E. EnablePARAMETER_SNIFFINGin DB1

Answer: AC

Explanation:
Last_Query_Plan_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in
sys.dm_exec_query_plan_stats.
Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT_QUERY_PROFILING database scoped configuration: ALTER DATABASE
SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure

NEW QUESTION 83
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.
You need to resolve the performance issues. What should you do?

A. Implement memory-optimized tables.
B. Run thedbcc flushprocindBcommand.
C. Replace the sequential index keys with nonsequential keys.
D. Run thedbcc dbreindexcommand.

Answer: A

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

NEW QUESTION 87

- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image.
You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance.
What should you do first?

A. From the SQL Server Installation Center on SQLVM1. run the Edition Upgrade wizard.
B. From SQLVM1, uninstall the SQL Server instance.
C. From the SQL Server Installation Center on SQLVM1. run the Repair wizard.
D. From the Azure portal, reconfigure SQLVM1.

Answer: B

NEW QUESTION 88
- (Exam Topic 5)
You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation. Which service should you use to ingest the data?

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

Answer: D

Explanation:
The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.
Reference:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions

NEW QUESTION 90
- (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 94
- (Exam Topic 5)
You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features. You plan to migrate to Azure SQL.
Which service should you use?

A. Azure SQL Database
B. SQL Server on an Azure Virtual Machine
C. Azure SQL Managed Instance
D. Azure Database for MySQL

Answer: B

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database-overview

NEW QUESTION 98
- (Exam Topic 5)
You have an Azure subscription that contains a logical SQL server named Serverl. The master database of Serverl contains a user named Userl. You need to ensure that Userl can create databases on Serverl. Which database role should you assign to Userl?

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

Answer: B

NEW QUESTION 101
- (Exam Topic 5)
You have an Azure SQL Database elastic pool that contains 10 databases. You receive the following alert.
Msg 1132, Level 16, State 1, Line 1
The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.
You need to resolve the alert. The solution must minimize administrative effort.
Which three actions can you perform? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Delete data from a database.
B. Remove a database from the pool.
C. Increase the maximum storage of the elastic pool.
D. Shrink individual databases.
E. Enable data compression.

Answer: BCD

NEW QUESTION 102
- (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.	○	○
DB1 will be placed in an offline state.	○	○
You are performing a tail-log backup.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The log file will be truncated.	○	○
DB1 will be placed in an offline state.	○	○
You are performing a tail-log backup.	○	○

NEW QUESTION 104
- (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? ⓘ

`14`

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:

| 52 | Week(s) ⌄ |

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

| 60 | Week(s) ⌄ |

Yearly LTR Backups
Keep an annual backup for:

| 10 | Year(s) ⌄ |

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

| Week 52 ⌄ |

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]

- 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

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 107
- (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 109
- (Exam Topic 5)
A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
(
    EmployeeID      INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName       VARCHAR(100) NOT NULL,
    LastName        VARCHAR(100) NOT NULL,
    Title           VARCHAR(100) NULL,
    LastHireDate    DATETIME NULL,
    StreetAddress1  VARCHAR(500) NOT NULL,
    StreetAddress2  VARCHAR(500) NOT NULL,
    StreetAddress3  VARCHAR(500) NOT NULL,
    City            VARCHAR(200) NOT NULL,
    StateName       VARCHAR(20) NOT NULL,
    Salary          VARCHAR(20) NULL,
    PhoneNumber     VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance.
Which two actions should you recommend? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

A. ChangeSalaryto themoneydata type.
B. ChangePhoneNumberto thefloatdata type.
C. ChangeLastHireDateto thedatetime2(7)data type.
D. ChangePhoneNumberto thebigintdata type.
E. ChangeLastHireDateto thedatedata type.

Answer: AE

NEW QUESTION 110
- (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 112
- (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 114
- (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:
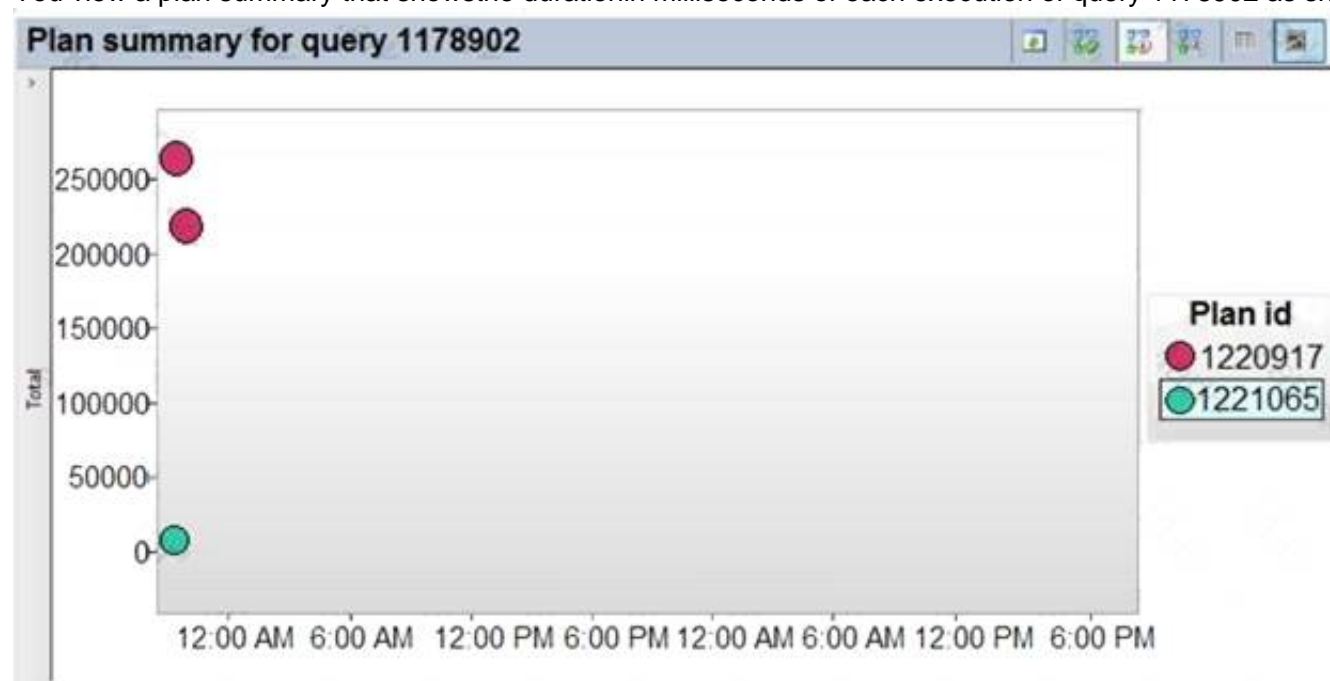https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet

NEW QUESTION 118
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1.
You view a plan summary that showsthe durationin 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 theDBCC FREEPROCCACHEcommand.
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 121
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 124
- (Exam Topic 5)
You are building an Azure Stream Analytics job to retrieve game data.
You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.
How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
SELECT _____ ▼ as HighestScore
        Collect(Score)
        CollectTop(1)OVER(ORDER BY Score Desc)
        Game, MAX(Score)
        TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

 FROM input TIMESTAMP BY CreatedAt

 GROUP BY _____ ▼
          Game
          Hopping(minute, 5)
          Tumbling(minute, 5)
          Windows(TumblingWindow(minute, 5), Hopping(minute, 5))
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)
TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.
Analytic Function Syntax:
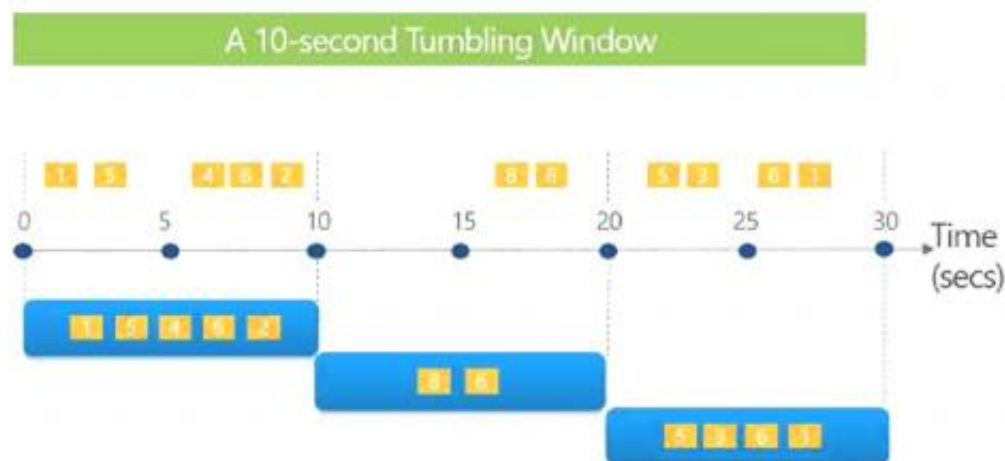TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT
DURATION clause> [<WHEN clause>])
Box 2: Tumbling(minute 5)
Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

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

NEW QUESTION 127
- (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 MiI.
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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Backup type: `Full only` ▼
- **Full only**
- Full and differential
- Full and transaction log
- Transaction log only

To store the backups: `Append each backup to a separate file.` ▼
- Append all the backups to a single file.
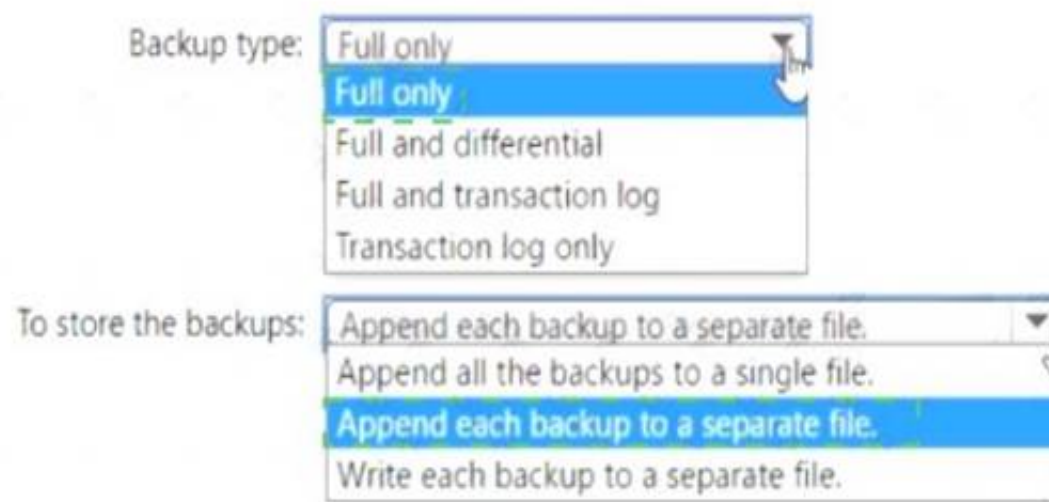- **Append each backup to a separate file.**
- Write each backup to a separate file.

NEW QUESTION 131
- (Exam Topic 5)
You create five Azure SQL Database instances on the same logical server.
In each database, you create a user for an Azure Active Directory (Azure AD) user named User1. User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.
You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.
What should you do?

A. Create User1 in the master database.
B. Assign User1 the db_datareader role for the master database.
C. Assign User1 the db_datareader role for the databases that UserI creates.
D. Grantselecton sys.databases to public in the master database.

Answer: A

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

NEW QUESTION 133
- (Exam Topic 5)
You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.
You execute the Transact-SQL query shown in the following exhibit.

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

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

The query results include only **[answer choice]** in the csv/busfare folder.
▼
- CSV files in the tripdata_2020 subfolder
- files that have files names beginning with "tripdata_2020"
- CSV files that have file names containing "tripdata_202"
- CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.
▼
- a header
- a data
- an empty

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 138
- (Exam Topic 5)
You have a SQL pool in Azure Synapse that contains a table named dbo.Customers. The table contains a column name Email.
You need to prevent nonadministrative users from seeing the full email addresses in the Email column. The users must see values in a format of aXXX@XXXX.com instead.
What should you do?

A. From the Azure portal, set a mask on the Email column.
B. From the Azure portal, set a sensitivity classification of Confidential for the Email column.
C. From Microsoft SQL Server Management Studio, set an email mask on the Email column.
D. From Microsoft SQL Server Management Studio, grant the SELECT permission to the users for all the columns in the dbo.Customers table except Email.

Answer: B

Explanation:
The Email masking method, which exposes the first letter and replaces the domain with XXX.com using a constant string prefix in the form of an email address.
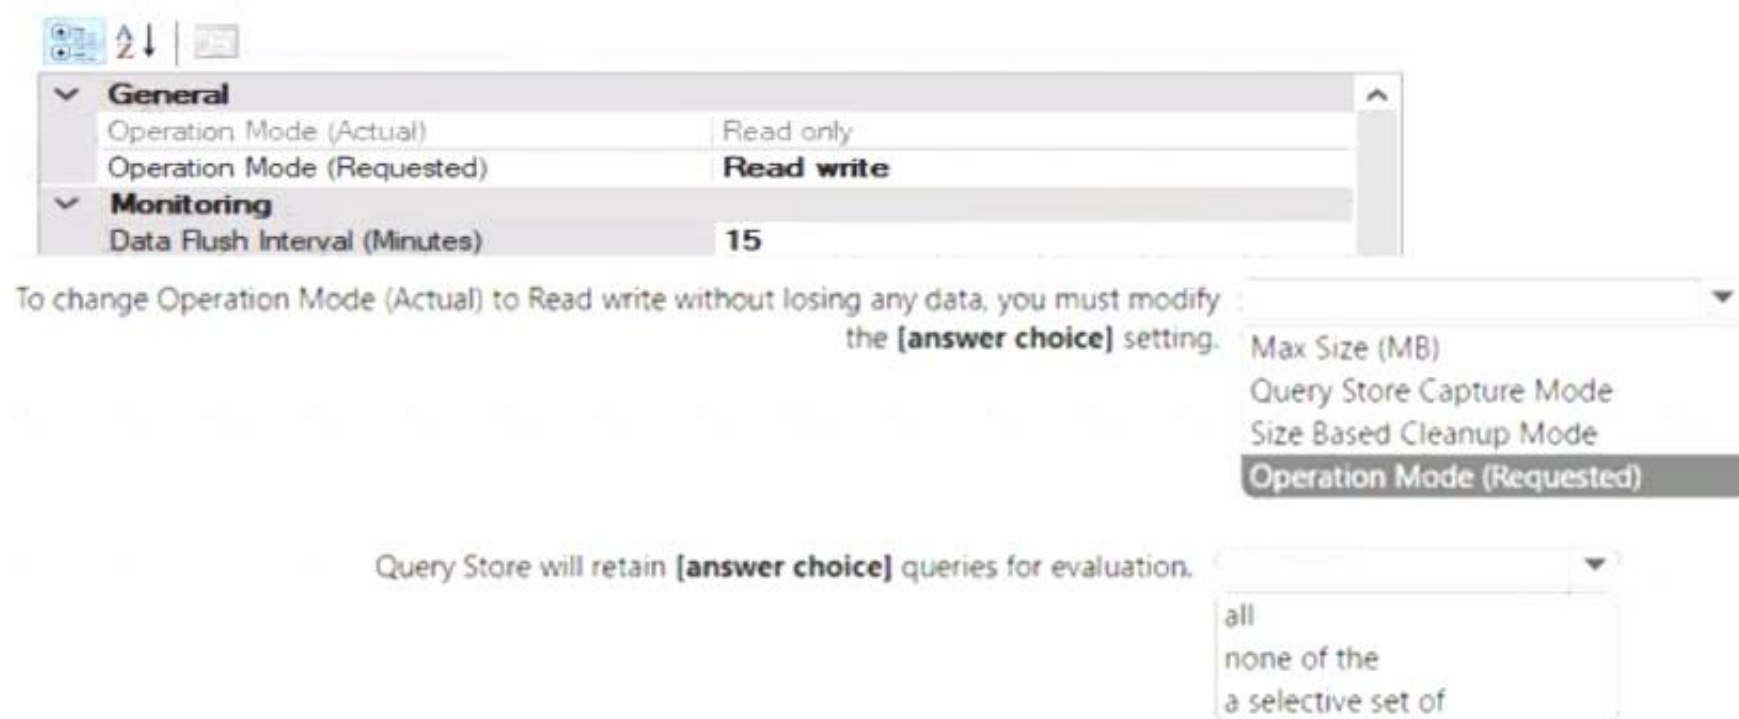Example: aXX@XXXX.com

NEW QUESTION 141
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text Description automatically generated

NEW QUESTION 146
- (Exam Topic 5)
You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.
Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted. You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the
amount of time required to delete old data.
How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.
You may need to drag the split bar between panes or scroll to view content.
NOTE:Each correct selection is worth one point.

Values

CustomerKey
HASH
ROUND_ROBIN
REPLICATE
OrderDateKey
SalesOrderNumber

Answer Area

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: HASH
Box 2: OrderDateKey
In most cases, table partitions are created on a date column.
A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management. For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool

NEW QUESTION 148
- (Exam Topic 5)
You have an Azure Synapse Analytics Apache Spark pool named Pool1.
You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.
You need to load the files into the tables. The solution must maintain the source data types. What should you do?

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

Answer: B

Explanation:
Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.
Serverless SQL pool enables you to query data in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.
To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.
The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-data-storage

NEW QUESTION 149
- (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: ❶

| Server | Azure defaults | Don't inherit |

❶ The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options ❶

	OPTION	DESIRED STATE	CURRENT STATE
	FORCE PLAN	ON / OFF / **INHERIT**	ON Auto-configured by Azure
	CREATE INDEX	ON / OFF / **INHERIT**	ON Auto-configured by Azure
	DROP INDEX	**ON** / OFF / 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.	○	○
Columns will be added to existing indexes automatically.	○	○
The query execution plan will revert to a previous plan if query performance degrades.	○	○

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 151
- (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 155
- (Exam Topic 5)
You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group.
You need to recommend a deployment solution that meets the following requirements:
• Provides a Service Level Agreement (SLA) of at least 99.95%
• Replicates databases in the same group synchronously
• Minimizes the latency of database writes What should you recommend?

A. Create a proximity group and an availability se
B. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group.
C. Create two proximity groups and a single availability se
D. Deploy both virtual machines to the availability se
E. Add one virtual machine to each proximity group.
F. Create two proximity groups and two availability set
G. Deploy each virtual machine to a unique availability se
H. Add one virtual machine to each proximity group.
I. Create a proximity group and two availability set
J. Deploy each virtual machine to a unique availability se
K. Add both virtual machines to the proximity group.

Answer: A

NEW QUESTION 157
- (Exam Topic 5)
You are developing an application that uses Azure Data Lake Storage Gen 2.
You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

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

Answer: C

Explanation:
A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:
What resources the client may access.
What permissions they have to those resources. How long the SAS is valid.
Note: Data Lake Storage Gen2 supports the following authorization mechanisms:
≫ Shared Key authorization
≫ Shared access signature (SAS) authorization
≫ Role-based access control (Azure RBAC)
≫ Shared Key authorization
≫ Shared access signature (SAS) authorization
≫ Role-based access control (Azure RBAC)
≫ Access control lists (ACL)
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

NEW QUESTION 159
- (Exam Topic 5)
You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.
You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication. You need to ensure that App1 can access DB2. The solution must meet the following requirements:
≫ App1 must be able to view only DB2.
≫ Administrative effort must be minimized. What should you create?

A. a contained database user for App1 on DB2
B. a login for App1 on Server1
C. a contained database user from an external provider for App1 on DB2
D. a contained database user from a Windows login for App1 on DB2

Answer: D

Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-databa

NEW QUESTION 163
- (Exam Topic 5)
You have an Azure SQL logical server. You run the following script.

```
CREATE DATABASE Sales
GO
CREATE TABLE [dbo].[Orders]
  (
    [OrderID] INT NOT NULL,
    [OrderDescription] NVARCHAR (MAX) NOT NULL,
    [Timestamp] Datetime2 NOT NULL
  )
  WITH (
    SYSTEM_VERSIONING = ON,
    LEDGER = ON
    );
GO
```

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 Orders table will allow only rows to be inserted.	○	○
To create additional tables in the Sales database, the LEDGER = ON parameter must be used.	○	○
To ensure that a timestamp is added to each row in the Orders table, the GENERATED ALWAYS	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Yes
No No

NEW QUESTION 165
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.
➢ Clustered columnstore indexes
➢ Automatic tuning
➢ Change tracking
➢ PolyBase
You plan to migrate DB1 to an Azure SQL database.
What feature should be removed or replaced before DB1 can be migrated?

A. Clustered columnstore indexes
B. PolyBase
C. Change tracking
D. Automatic tuning

Answer: B

Explanation:
This table lists the key features for PolyBase and the products in which they're available.

Feature	SQL Server (Beginning with 2016)	Azure SQL Database	Azure Synapse Analytics	Parallel Data Warehouse
Query Hadoop data with Transact-SQL	Yes	No	No	Yes
Import data from Hadoop	Yes	No	No	Yes
Export data to Hadoop	Yes	No	No	Yes
Query, import from, export to Azure HDInsight	No	No	No	No
Push down query computations to Hadoop	Yes	No	No	Yes
Import data from Azure Blob storage	Yes	Yes*	Yes	Yes
Export data to Azure Blob storage	Yes	No	Yes	Yes
Import data from Azure Data Lake Store	No	No	Yes	No
Export data to Azure Data Lake Store	No	No	Yes	No
Run PolyBase queries from Microsoft BI tools	Yes	No	Yes	Yes

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary

NEW QUESTION 167
- (Exam Topic 5)
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

Actions

- a virtual disk that uses the stripe layout
- a virtual disk that uses the mirror layout
- a volume
- a virtual disk that uses the simple layout
- a storage pool

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Follow these same steps to create striped virtual disk:
➤ Create Log Storage Pool.
➤ Create Virtual Disk
➤ Create Volume
Box 1: a storage pool
Box 2: a virtual disk that uses stripe layout
Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.
Box 3: a volume Reference:
https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/

NEW QUESTION 168

- (Exam Topic 5)
Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1. You have an Azure subscription.
You plan to migrate db1 to an Azure SQL managed instance.
You need to create the SQL managed instance. The solution must minimize the disk latency of the instance. Which service tier should you use?

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

Answer: D

NEW QUESTION 169
- (Exam Topic 5)
You have an Azure Data Factory pipeline that is triggered hourly. The pipeline has had 100% success for the past seven days.
The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

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

What is a possible cause of the error?

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

Answer: B

NEW QUESTION 172
- (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 176
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1. You use Query Performance Insight to monitor db1.
You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.
Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Configuration setting:

```
DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE
```

Value:

```
1
60
CUSTOM
ON
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated

NEW QUESTION 181
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. The database reports a CHECKSUM error.
You need to recover the database.
How should you complete the statements? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

USE master;

ALTER DATABASE [DB1] SET ▼ WITH ROLLBACK IMMEDIATE;

GO

```
OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY
```

DBCC CHECKDB ('DB1', ▼ WITH NO_INFOMSGS;

GO

```
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
```

ALTER DATABASE [DB1] SET ▼

GO

```
MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SINGLE_USER
The specified database must be in single-user mode to use one of the following repair options. Box 2: REPAIR_ALLOW_DATA_LOSS
REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.
Note: The REPAIR_ALLOW_DATA_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA_LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql

NEW QUESTION 186
- (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 188
- (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 192
- (Exam Topic 5)
You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.
You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

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

Answer: C

Explanation:
You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.
Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.
Reference:
https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html

NEW QUESTION 197
- (Exam Topic 5)
You are creating a managed data warehouse solution on Microsoft Azure.
You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.
You need to configure Azure Synapse Analytics to receive the data.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, chat or text message Description automatically generated
To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table.
Step 1: Create a master key on database.
* 1. Create a master key on the database. The master key is required to encrypt the credential secret. (Create a database scoped credential for Azure blob storage.)
Step 2: Create an external data source for Azure Blob storage.
* 2. Create an external data source with CREATE EXTERNAL DATA SOURCE.. Step 3: Create an external file format to map the parquet files.
* 3. Create an external file format with CREATE EXTERNAL FILE FORMAT. Step 4. Create an external table FactSalesOrderDetails
* 4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE. Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage

NEW QUESTION 201
- (Exam Topic 5)
You have an Azure SQL managed instance that hosts multiple databases.
You need to configure alerts for each database based on the diagnostics telemetry of the database. What should you use?

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

Answer: D

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

NEW QUESTION 206
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database. The database fails to respond to queries in a timely manner.
You need to identify whether the issue relates to resource_semaphore waits.
How should you complete the Transact-SQL query? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
SELECT    [ is_user_process ]  ▼
            [ wait_time ]
            [ wait_type ]

    SUM(wait_time) AS total_wait_time_ms

FROM sys.  [ dm_exec_query_stats ]  ▼
            [ dm_exec_requests ]
            [ query_store_query ]

    JOIN sys.dm_exec_sessions AS dmv2
        ON dmv1.session_id = dmv2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs

NEW QUESTION 211
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 instance named SQL1 that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named MI1 and an Azure Storage account named storage1.
You need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:

* Use block blob storage.
* Maximize security.

A. Generate a shared access signature (SAS)
B. Enable infrastructure encryption.
C. Create an access policy.
D. Rotate the storage keys

Answer: B

NEW QUESTION 213
- (Exam Topic 5)
You have an Azure SQL database named sqldb1.
You need to minimize the possibility of Query Store transitioning to a read-only state. What should you do?

A. Halve the value of Data Flush Interval.
B. Double the value of Statistics Collection Interval.
C. Halve the value of Statistics Collection Interval.
D. Double the value of Data Flush Interval.

Answer: A

Explanation:
The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval
(Minutes) option. If Query Store has breached the
Maximum size limit between storage size checks, it transitions to read-only mode. Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store

NEW QUESTION 216
- (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 218
- (Exam Topic 5)
You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:
⟩ Support failovers that do not require client applications to change their connection strings.
⟩ Replicate the database to a secondary Azure region.
⟩ Support failover to the secondary region. What should you include in the recommendation?

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

Answer: A

Explanation:
Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same
or different data center (region).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview

NEW QUESTION 219
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy an Azure SQL managed instance that meets the following requirements:
•Optimize latency.
•Maximize the memory-to-vCore ratio.
Which service tier and hardware generation should you use? To answer, select the apocopate 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 224
- (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.

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 227
- (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


```
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 228
- (Exam Topic 2)
Which audit log destination should you use to meet the monitoring requirements?

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

Answer: C

Explanation:
Scenario: Use a single dashboard to review security and audit data for all the PaaS databases.
With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics.
Note: Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/visualize/tutorial-logs-dashboards

NEW QUESTION 229
- (Exam Topic 1)
You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance requirements.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

| Create a user in the master database |

| Modify the Azure SQL server administrator account |

| Create contained database users |

| Create an Azure AD administrator for the logical server |

| Connect to the databases by using an Azure AD account |

| Enable the contained database authentication option |

Answer Area

A. Mastered
B. Not Mastered

Answer: A

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

≫ Create and populate Azure AD.

≫ Optional: Associate or change the active directory that is currently associated with your Azure Subscription.

≫ Create an Azure Active Directory administrator. (Step 1)

≫ Configure your client computers.

≫ Create contained database users in your database mapped to Azure AD identities. (Step 2)

≫ Connect to your database by using Azure AD identities. (Step 3)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-overview

NEW QUESTION 231
- (Exam Topic 1)
You need to implement authentication for ResearchDB1. The solution must meet the security and compliance requirements.
What should you run as part of the implementation?

A. CREATE LOGINand theFROM WINDOWSclause
B. CREATE USERand theFROM CERTIFICATEclause
C. CREATE USERand theFROM LOGINclause
D. CREATE USERand theASYMMETRIC KEYclause
E. CREATE USERand theFROM EXTERNAL PROVIDERclause

Answer: E

Explanation:
Scenario: Authenticate database users by using Active Directory credentials.
(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.) Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user: CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql

NEW QUESTION 234
- (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 239
- (Exam Topic 1)
You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.
What should you include in the plan?

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

Answer: D

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

NEW QUESTION 244
......

Relate Links

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

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

Contact us

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

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

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

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