

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 SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1.

You need to test impersonation of user1 in db1 by running aSELECTstatement and returning to the original execution context.

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

NOTE:Each correct selection is worth one point.

```
EXECUTE AS  = 'user1@contoso.com'
```

CALLER

LOGIN

OWNER

USER

GO

```
SELECT SUSER_SNAME ()
```

REVERT

REVOKE

ROLLBACK

GO

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15> <https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-sname-transact-sql?view=sql-server-ver15>

NEW QUESTION 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.

- 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 Azure subscription.

You need to deploy an Azure SQL managed instance by using an Azure Resource Manager (ARM) template. The solution must meet the following requirements:

The SQL managed instance must be assigned a unique identity.

The SQL managed instance must be available in the event of an Azure datacenter outage.

How should you complete the template? 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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
{
  "type": "Microsoft.Sql/managedInstances",
  -
},
"identity": {
  "type": "SystemAssigned"
},
"dependsOn": [
  "[parameters('virtualNetworkName')]"
],
"properties": {
  "administratorLogin": "[parameters('administratorLogin')]",
  "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
  "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
  "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
  parameters('virtualNetworkName'), parameters('subnetName'))]", "storageSizeInGB": 8192,
  "vCores": 80, "licenseType": "BasePrice",
  "zoneRedundant": "True"
}
}
```

NEW QUESTION 5

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

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

NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 9

- (Exam Topic 5)

You 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 10

- (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 SQL database named DB 1 in the General Purpose service tier. You need to monitor DB 1 by using SQL Insights. What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

To collect monitoring data, use:

- 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:

Box 1 = Azure Monitor Agent Box 2 = An Azure SQL database
<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql>

NEW QUESTION 12

- (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 SEC
```

- 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_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_SINGLE_EVENT_LOSS
NO_EVENT_LOSS
MAX_DISPATCH_LATENCY = 3 SE
```

NEW QUESTION 14

- (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:

Built-in role:

- 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 16

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1.

You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

Answer Area

Backup type:

▼
Full and log backups only
Full backup only
Log backup only

Backup option:

▼
WITH CHECKSUM
WITH NOINIT
WITH UNLOAD

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Full and log backups only

Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn't support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.

Box 2: WITH CHECKSUM

Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database Migration Service only supports backups created using checksum.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online>

NEW QUESTION 21

- (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 26

- (Exam Topic 5)

You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1.

You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.

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

Actions

```
Run CREATE LOGIN user1@contoso.com FROM EXTERNAL PROVIDER on the master database.
```

```
Run ALTER SERVER ROLE securityadmin ADD MEMBER user1@contoso.com.
```

```
Create a managed identity for SQLMI1.
```

```
Grant SQLMI1 read access to Azure AD.
```

```
Run CREATE USER user1@contoso.com FROM LOGIN user1@contoso.com.
```

Answer Area

Navigation icons: Left arrow, Right arrow, Up arrow, Down arrow.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated with medium confidence

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/aad-security-configure-tutorial>

NEW QUESTION 29

- (Exam Topic 5)

You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:

- > ProductID
- > ItemPrice
- > LineTotal
- > Quantity
- > StoreID
- > Minute
- > Month
- > Hour
- > Year
- > Day

You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.

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

NOTE: Each correct selection is worth one point.

```
df.write
```

<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> .bucketBy	<input type="checkbox"/> ("*")
<input type="checkbox"/> .partitionBy	<input type="checkbox"/> ("StoreID", "Hour")
<input type="checkbox"/> .range	<input type="checkbox"/> ("StoreID", "Year", "Month", "Day", "Hour")
<input type="checkbox"/> .sortBy	<input type="checkbox"/> ("Year", "Month", "Day", "Hour", "StoreID")

```
.mode("append")
```

<input type="checkbox"/>
<input type="checkbox"/> .csv("/Purchases")
<input type="checkbox"/> .json("/Purchases")
<input type="checkbox"/> .parquet("/Purchases")
<input type="checkbox"/> .saveAsTable("/Purchases")

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: .partitionBy Example:
df.write.partitionBy("y","m","d") mode(SaveMode.Append) parquet("/data/hive/warehouse/db_name.db/" + tableName) Box 2:
("Year","Month","Day","Hour","StoreID")
Box 3: .parquet("/Purchases") Reference:
<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partiti>

NEW QUESTION 31

- (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 32

- (Exam Topic 5)

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

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

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

Create: ▼

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

Control access to the rows by using: ▼

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Create: ▼

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

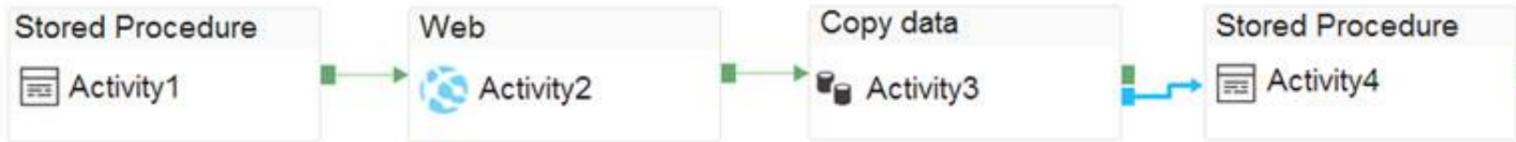
Control access to the rows by using: ▼

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

NEW QUESTION 33

- (Exam Topic 5)

You have an Azure data factory that has two pipelines named PipelineA and PipelineB. PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



You create an alert for the data factory that uses Failed pipeline runs metrics for both pipelines and all failure types. The metric has the following settings:

- > Operator: Greater than
- > Aggregation type: Total
- > Threshold value: 2
- > Aggregation granularity (Period): 5 minutes
- > Frequency of evaluation: Every 5 minutes

Data Factory monitoring records the failures shown in the following table.

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

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

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Box 1: No

Just one failure within the 5-minute interval. Box 2: No

Just two failures within the 5-minute interval. Box 3: No

Just two failures within the 5-minute interval. Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview>

NEW QUESTION 36

- (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 37

- (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 41

- (Exam Topic 5)

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

You need to ensure that User1 can create a scheduled task to perform a full backup of DB1. The solution must use the principle of least privilege.

Which built-in database role should you assign to User1?

- A. SQLAgentReaderRole
- B. db.owner
- C. SQLAgentOperatorRole
- D. SQLAgentUserRole

Answer: C

NEW QUESTION 45

- (Exam Topic 5)

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

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

The records contain two applicants at most.

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

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

NOTE:Each correct selection is worth one point.

▼ applications

CREATE EXTERNAL TABLE

CREATE TABLE

CREATE VIEW

```
WITH (
  LOCATION = 'applications/',
  DATA_SOURCE = applications_ds,
  FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
  (BULK 'https://contoso1.dfs.core.windows.net/
  applications/year=*/*.parquet',
  FORMAT = 'PARQUET') AS [r]
GO
```

CROSS APPLY

OPENJSON

OPENROWSET

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: CREATE EXTERNAL TABLE

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

Syntax:

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

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

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

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

Example: AS

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

```
OPENROWSET(BULK
```

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

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

```
GROUP BY decennialTime, stateName GO
```

Reference:

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

NEW QUESTION 50

- (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 52

- (Exam Topic 5)

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

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

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance. Solution: You create additional tempdb files. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Reference:

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

NEW QUESTION 54

- (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 `nvarchar(MAX)`.
- C. Change the column type to `varchar(MAX)`.
- 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 56

- (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 60

- (Exam Topic 5)

You have an Azure SQL managed instance.

You need to restore a database named DB1 by using Transact-SQL.

Which command should you run? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

RESTORE DB1 FROM

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

NEW QUESTION 65

- (Exam Topic 5)

You have a version-8.0 Azure Database for MySQL database.

You need to identify which database queries consume the most resources. Which tool should you use?

- A. Query Store
- B. Metrics
- C. Query Performance Insight
- D. Alerts

Answer: A

Explanation:

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time. Query Store simplifies performance troubleshooting by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns. Data for all users, databases, and queries is stored in the mysql schema database in the Azure Database for MySQL instance. Reference: <https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store>

NEW QUESTION 66

- (Exam Topic 5)

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

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

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

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

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

⬅

➡

⬆

⬇

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

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

NEW QUESTION 70

- (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 74

- (Exam Topic 5)

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL

You need to use PowerShell to enable and configure automated patching for SQL The solution must include both SQL Server and Windows security updates.

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

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2
-MaintenanceWindowDuration 120 -PatchCategory "Important"
```

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
```



```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

```
-SQLManagementType
```



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2
-MaintenanceWindowDuration 120 -PatchCategory "Important"
```

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
```



```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

```
-SQLManagementType
```



NEW QUESTION 78

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machines.

You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs. The solution must use the principle of least privilege.

Which role should you assign to User1?

- A. sysadmin
- B. SQLAgentUserRole
- C. SQLAgentReaderRole
- D. SQLAgentOperatorRole

Answer: A

NEW QUESTION 79

- (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 81

- (Exam Topic 5)

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

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

Answer: B

Explanation:

A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.

Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key>

NEW QUESTION 86

- (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 endpoint
- D. an ExpressRoute gateway

Answer: C

Explanation:

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.

The service could be an Azure service such as:

- > Azure Storage
- > Azure Cosmos DB
- > Azure SQL Database
- > Your own service using a Private Link Service. Reference:

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

NEW QUESTION 89

- (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 db1.
- C. Decrease the auto-pause delay for db1.
- D. Enable automatic tuning for db1.

Answer: D

NEW QUESTION 90

- (Exam Topic 5)

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 95

- (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 98

- (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 100

- (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 103

- (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 <code>CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;</code>	
On DatabaseA, run <code>CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';</code>	
On DatabaseA, run <code>ALTER ROLE db_datareader ADD MEMBER [App1];</code>	<div style="display: flex; justify-content: space-around;"> ⏪ ⏩ </div>
On the master database, run <code>CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';</code>	<div style="display: flex; justify-content: space-around;"> ⏩ ⏪ </div>
On DatabaseA, run <code>CREATE USER [App1] FROM LOGIN [App1];</code>	

- 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 104

- (Exam Topic 5)

You have an Azure SQL database named DB1.

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

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Hyperscale
- B. Azure SQL Database managed instance General Purpose
- C. Azure SQL Database Premium
- D. Azure SQL Database Basic

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 105

- (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 109

- (Exam Topic 5)

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

You need to record whenever users query the CustomerPII table.

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

NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:

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

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

Note:

The Server Audit Specification object belongs to an audit.

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

Reference:

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

NEW QUESTION 114

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL database named SQL1. SQL1 is in an Azure region that does not support availability zones.

You need to ensure that you have a secondary replica of SQL1 in the same region. What should you use?

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

Answer: C

NEW QUESTION 115

- (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

[app1]

CREDENTIAL
LOGIN
USER

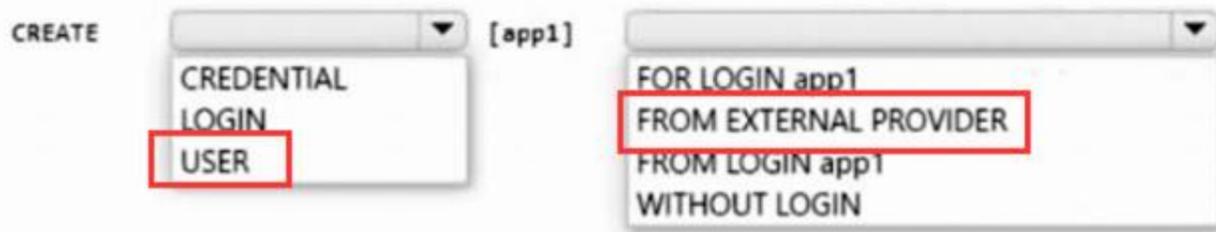
FOR LOGIN app1
FROM EXTERNAL PROVIDER
FROM LOGIN app1
WITHOUT LOGIN

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 116

- (Exam Topic 5)

You have an Azure SQL managed instance.

You need to enable SQL Agent Job email notifications. What should you do?

- A. Use the Agent XPs option.
- B. Enable the SQL Server Agent.
- C. Run the sp_configure command.
- D. Run the sp_set_agent_properties command.

Answer: C

NEW QUESTION 120

- (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 122

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. The solution must meet the following requirements:

- Minimize how long it takes to fail over.
- Maintain existing connections to the primary replica during a failover. What should you do?

- A. Connect each virtual machine to a single subnet on a single virtual network.
- B. Connect each virtual machine to a single subnet on a virtual network
- C. Deploy a standard Azure load balancer.
- D. Connect each virtual machine to a different subnet on a single virtual network.
- E. Connect each virtual machine to a different subnet on a virtual network
- F. Deploy a basic Azure load balancer.

Answer: C

NEW QUESTION 126

- (Exam Topic 5)

You have an Azure SQL database.

You need to identify whether a delayed query execution is associated to a RESOURCE wait.

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

NOTE: Each correct selection is worth one point.

Answer Area

```

SELECT [wait_type]
       [wait_type]
       [context_info]
       [wait_resource]

SUM(wait_time) AS total_wait_time_ms

FROM [sys.dm_exec_requests] AS dev1
JOIN [sys.dm_exec_requests]
     [sys.dm_exec_connections]
     [sys.dm_db_partition_stats] AS dev2 ON dev1.session_id = dev2.session_id

WHERE is_user_process = 1

GROUP BY TARGET1

ORDER BY SUM(wait_time) DESC;
    
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

SELECT [wait_type]
       [wait_type]
       [context_info]
       [wait_resource]

SUM(wait_time) AS total_wait_time_ms

FROM [sys.dm_exec_requests] AS dev1
JOIN [sys.dm_exec_requests]
     [sys.dm_exec_connections]
     [sys.dm_db_partition_stats] AS dev2 ON dev1.session_id = dev2.session_id

WHERE is_user_process = 1

GROUP BY TARGET1

ORDER BY SUM(wait_time) DESC;
    
```

NEW QUESTION 128

- (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 130

- (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 134

- (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 137

- (Exam Topic 5)

You have an Azure SQL managed instance named MI1.

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

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

Answer: B

NEW QUESTION 138

- (Exam Topic 5)

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

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

What should you include in the recommendation?

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

Answer: B

Explanation:

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

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

Example:

XXXX-XXXX-XXXX-1234

NEW QUESTION 141

- (Exam Topic 5)

You have an Azure subscription that contains two Azure SQL managed instances named SQLMI1 and SQLMI2 . SQLMI2 contains a database named DB1 and a user named User1.

User1 drops DB1.

You need to perform a point-in-time restore of DB1 to SQLMI2.

- A. Azure CLI
- B. Transact-SQL
- C. The Azure portal
- D. Azure PowerShell

Answer: C

NEW QUESTION 146

- (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

Answer Area

- 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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Answer Area

- 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.



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

NEW QUESTION 151

- (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 153

- (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 155

- (Exam Topic 5)

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection. You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database. What should you use?

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

Answer: C

Explanation:

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

NEW QUESTION 157

- (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 159

- (Exam Topic 5)

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

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

You have an Azure 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 162

- (Exam Topic 5)

You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

Name	Type
Column1	Ntext
Column2	Geometry
Column3	Image
Column4	Varchar
Column5	Datetime2

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point

- A. Column1
- B. Column2
- C. Column3
- D. Column4
- E. Column5

Answer: AD

NEW QUESTION 163

- (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 166

- (Exam Topic 5)

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

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

What should you use?

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

Answer: A

Explanation:

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage

demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.

Reference:

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

NEW QUESTION 170

- (Exam Topic 5)

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

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

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

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

Answer: BD

Explanation:

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

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

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

Reference:

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

NEW QUESTION 174

- (Exam Topic 5)

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

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

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

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 178

- (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 you do?

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

Answer: D

NEW QUESTION 179

- (Exam Topic 5)

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

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

You have 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:

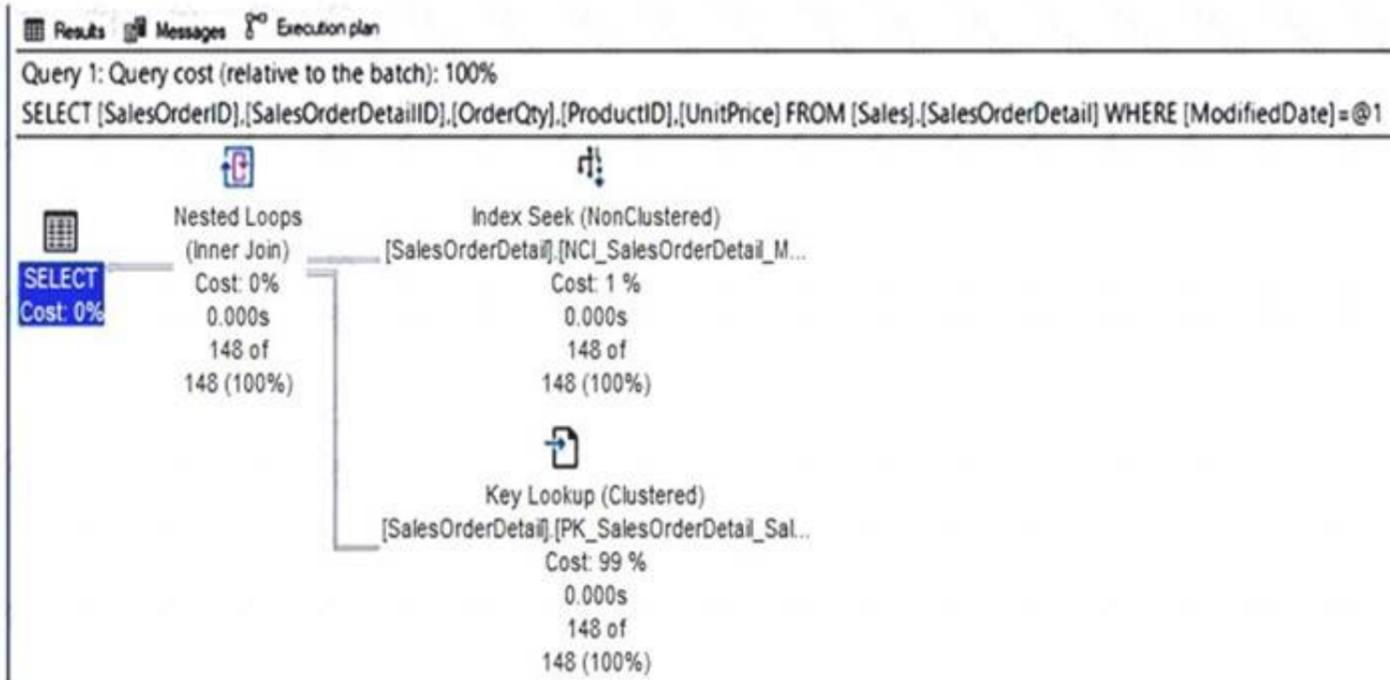
<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION 180

- (Exam Topic 5)

You have an Azure SQL database.

You have a query and the associated execution plan 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 performance issue stems from the [answer choice] operator.

▼

Select
 Index Seek
 Key Lookup
 Nested Loops

The performance issue can be resolved by adding include columns to the [answer choice].

▼

heap
 clustered index
 nonclustered index

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: Key Lookup

The Key Lookup cost is 99% so that is the performance bottleneck. Box 2: nonclustered index

The key lookup on the clustered index is used because the nonclustered index does not include the required columns to resolve the query. If you add the required columns to the nonclustered index, the key lookup will not be required.

NEW QUESTION 181

- (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 184

- (Exam Topic 5)

You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1.

You plan to create a database named DB1 in Pool1.

You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.

Which format should you use for the tables in DB1?

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

Answer: C

Explanation:

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

For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.

Reference:

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

NEW QUESTION 187

- (Exam Topic 5)

You create five Azure SQL Database instances on the same logical server.

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1. User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.

You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.

What should you do?

- A. Create User1 in the master database.
- B. Assign User1 the db_datareader role for the master database.
- C. Assign User1 the db_datareader role for the databases that User1 creates.
- D. 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 191

- (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 195

- (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 196

- (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 at](#)

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 at](#)

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 197

- (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 200

- (Exam Topic 5)

You have an Azure subscription.
 You plan to deploy an Azure SQL database by using an Azure Resource Manager template.
 How should you complete the template? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

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

NEW QUESTION 201

- (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.

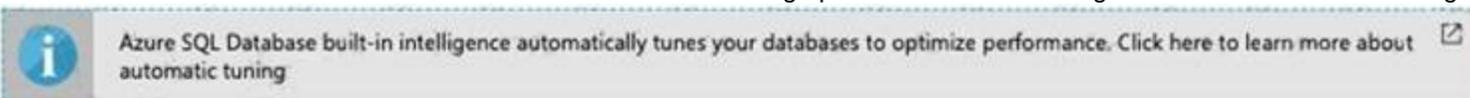
Reference:

<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 205

- (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.



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 209

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.

You need to create a data source in Pool1 that you can reference when you create the external table. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

CREATE EXTERNAL DATA SOURCE source1

WITH

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

)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, table Description automatically generated

Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2 CREATE EXTERNAL DATA SOURCE YellowTaxi WITH (LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/', TYPE = HADOOP)

Box 2: HADOOP

Reference:

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

NEW QUESTION 213

- (Exam Topic 5)

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

SQL1 contains an Extended Events session named session1 that captures Microsoft SQL Server events. You need to correlate the session events with events captured by Event Tracing for Windows (ETW). What should you do for session1?

- A. Modify the Set Session Event Filters settings.
- B. Add a target.
- C. Add an action.
- D. Modify the Specify Session Data Storage settings.

Answer: B

NEW QUESTION 215

- (Exam Topic 5)

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

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

You have an Azure SQL database named Sales.

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

- > During normal operations, provide at least two readable copies of Sales.
- > Ensure that Sales remains available if a datacenter fails.

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 216

- (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 221

- (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	<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: 40px; height: 40px; display: flex; align-items: center; justify-content: center;"> ⬆ </div> </div>
Create an Azure Policy initiative	<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: 40px; height: 40px; display: flex; align-items: center; justify-content: center;"> ⬇ </div> </div>
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 226

- (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 227

- (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 228

- (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 231

- (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 235

- (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 240

- (Exam Topic 5)

You are building an Azure virtual machine.

You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.

You need to configure storage for the virtual machine to support the database.

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

Actions	Answer Area
a virtual disk that uses the stripe layout	
a virtual disk that uses the mirror layout	
a volume	
a virtual disk that uses the simple layout	
a storage pool	




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Follow these same steps to create striped virtual disk:

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

Box 1: a storage pool

Box 2: a virtual disk that uses stripe layout

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

Box 3: a volume Reference:

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

NEW QUESTION 242

- (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 245

- (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 247

- (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 248

- (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 252

- (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 254

- (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 259

- (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 260

- (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 265

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL database. The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource_semaphore waits.

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

NOTE: Each correct selection is worth one point.

```

SELECT 
      
      
      SUM(wait_time) AS total_wait_time_ms
FROM sys.
      JOIN sys.dm_exec_sessions AS dmvs1
      ON dmvs1.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 270

- (Exam Topic 5)

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1. You need to encrypt the data in Col1.

Which four actions should you perform for DB1 in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- Create a database master key.
- Create a column master key.
- Open the symmetric key.
- Create a certificate.
- Update Col1.
- Create a symmetric key.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

Reference:

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

NEW QUESTION 272

- (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.

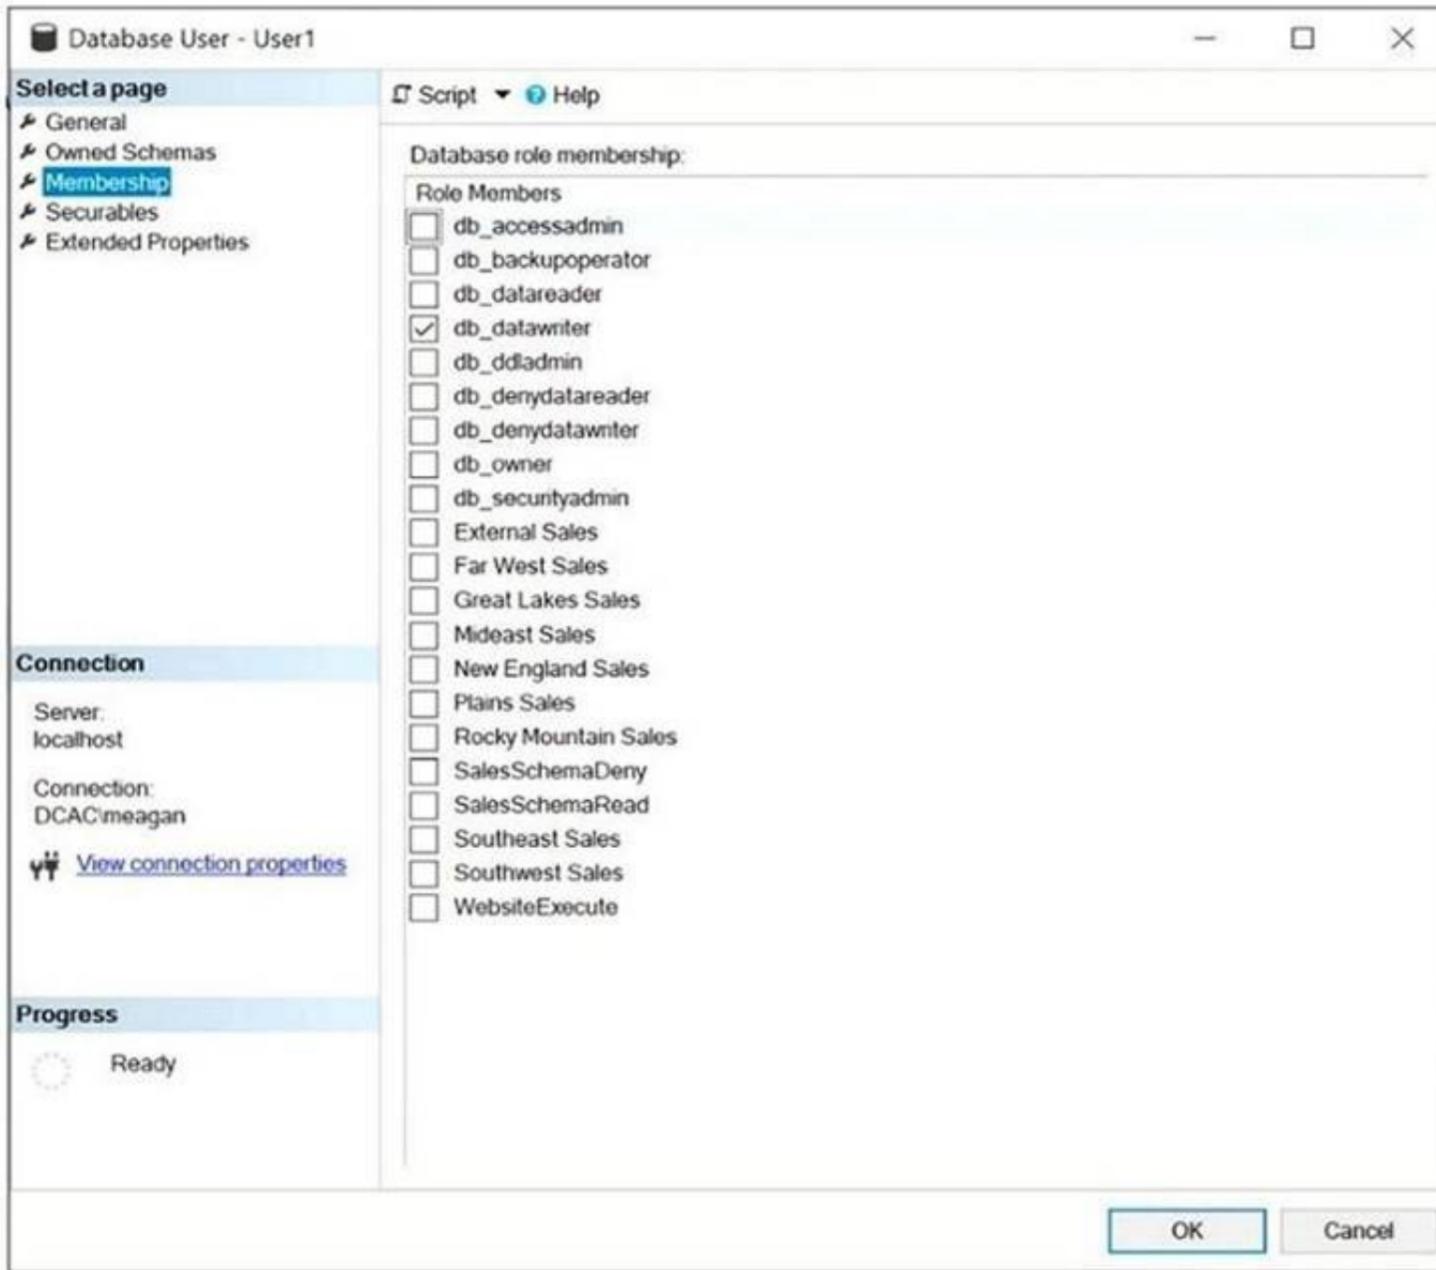
Reference:

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

NEW QUESTION 277

- (Exam Topic 5)

You have a Microsoft SQL Server database named DB1 that contains a table named Table1. The database role membership for a user named User1 is 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.

User1 can [answer choice].

	▼
add a column to Table1	
delete a row from Table1	
delete Table1	

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

	▼
db_datareader	
db_ddladmin	
db_denydatareader	
db_denydatawriter	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: delete a row from Table1
 Members of the db_datawriter fixed database role can add, delete, or change data in all user tables. Box 2: db_datareader
 Members of the db_datareader fixed database role can read all data from all user tables. Reference:
<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

NEW QUESTION 280

- (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 284

- (Exam Topic 5)

You have a 50-TB Microsoft SQL Server database named DB1.

You need to reduce the time it takes to perform database consistency checks of DB1.

Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

DBCC CHECKDB ([DB1],

	▼
NOINDEX	
REPAIR_FAST	
REPAIR_REBUILD	

 with

	▼
ALL_ERRORMSGs	
NO_INFOMSGs	
PHYSICAL_ONLY	

)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated with low confidence

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-ser>

NEW QUESTION 287

- (Exam Topic 5)

You have a burstable Azure virtual machine named VM1 that hosts an instance of Microsoft SQL Server. You need to attach an Azure ultra disk to VM1. The solution must minimize downtime on VM1.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Attach the ultra disk.
Stop and deallocate VM1.
Set Enable Ultra disk compatibility to Yes .
Resize VM1.
Start VM1.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Attach the ultra disk.
Stop and deallocate VM1.
Set Enable Ultra disk compatibility to Yes .
Resize VM1.
Start VM1.

Answer Area

Stop and deallocate VM1.
Attach the ultra disk.
Set Enable Ultra disk compatibility to Yes .
Resize VM1.
Start VM1.

NEW QUESTION 291

- (Exam Topic 5)

You have an on-premises app named App1 that stores data in an on-premises Microsoft SQL Server 2016 database named DB1.

You plan to deploy additional instances of App1 to separate Azure regions. Each region will have a separate instance of App1 and DB1. The separate instances of DB1 will sync by using Azure SQL Data Sync.

You need to recommend a database service for the deployment. The solution must minimize administrative effort.

What should you include in the recommendation?

- A. Azure SQL Managed instance
- B. Azure SQL Database single database
- C. Azure Database for PostgreSQL
- D. SQL Server on Azure virtual machines

Answer: B

Explanation:

Azure SQL Database single database supports Data Sync. Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

NEW QUESTION 295

- (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 300

- (Exam Topic 3)

Which counter should you monitor for real-time processing to meet the technical requirements?

- A. SU% Utilization
- B. CPU% utilization
- C. Concurrent users
- D. Data Conversion Errors

Answer: B

Explanation:

Scenario: Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

To monitor the performance of a database in Azure SQL Database and Azure SQL Managed Instance, start by monitoring the CPU and IO resources used by your workload relative to the level of database performance you chose in selecting a particular service tier and performance level.

Reference:

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

NEW QUESTION 303

- (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

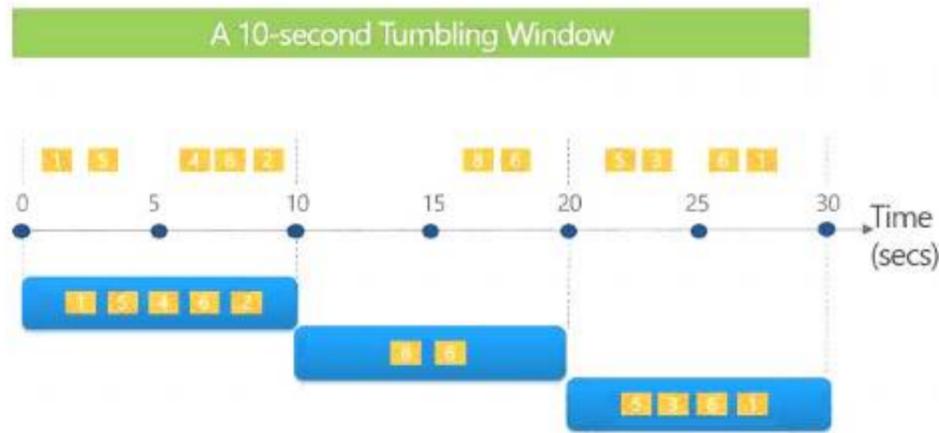
Explanation:

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 304

- (Exam Topic 2)
 Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

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

Answer: A

Explanation:

There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed. Reference:
<https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview>

NEW QUESTION 307

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

Purchasing model:

Deployment option:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: DTU
 Scenario:
 > The 30 new databases must scale automatically.
 > Once all requirements are met, minimize costs whenever possible.
 You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.
 In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.
 Box 2: An Azure SQL database elastic pool
 Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage

demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database. Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview> <https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview>

NEW QUESTION 310

- (Exam Topic 1)

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

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

Azure Database Migration Service pricing tier:

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

Required Azure resource:

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure Database Migration service Box 1: Premium 4-VCore

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

➤ Minimize downtime during the migration of the SERVER1 databases.

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

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

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

NEW QUESTION 314

- (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

- 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.

Answer Area

◀	▶
▶	◀

- 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 319

- (Exam Topic 1)

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements. How should you collect and stream metrics? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

Collect metrics from:

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

Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: The server, the elastic pool, and the database

Senario:

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

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

Box 2: Azure Event hubs

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

NEW QUESTION 321

.....

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](#)