

Microsoft.DP-300-KR.v2026-06-08.q157

□□□□:	DP-300-KR
□□□□:	Administering Relational Databases on Microsoft Azure (DP-300 Korean Version)
□□□:	Microsoft
□□ □□ □□□:	157
□□:	v2026-06-08
# □□ □:	109
# □□ □□□:	1570
https://www.krdump.com/Microsoft.DP-300-KR.v2026-06-08.q157.html	

NEW QUESTION: 1

Azure Stream Analytics □□□ □□□□.

□□□ □□□ □□□□ □□□ □□□□□□□ □□□ □□□□ □□□.

SU % Utilization □□□ □□□□□ □□□□□.

□□ □ □□ □□ □□□ □□□□□ □□□? □ □□□ □□□□ □□□ □□□□□.

□□: □□□ □□□ □□ 1□□ □□□ □□□□.

- A. □□ □□ □□□
- B. □□□ □□□ □□□
- C. □□□□ □□ □□□
- D. □□□□ □□
- E. □□ □□□

Answer: (SHOW ANSWER)

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

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

ference:

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

NEW QUESTION: 2

□□ □□ □□□ □□ 40□□ Azure SQL □□□□□□□ □□□□. □□ □□□□□□□ □□□ Azure SQL Database □□□ □□□□.

□ □□□ □□ □□□□□□□□ □□□□ □□□□ □ □□□ □□ □□□.

□□ □ □□ □□□ □□□□ □□□? □ □□□ □□□□ □□□ □□□□□.

□□: □□□ □□□ □□ 1□□ □□□ □□□□.

- A. □ □□ □□(RLS)□ □□□□□.
- B. □ □□□□□□□ □□□□ □□□□□.
- C. □□□□□□ □□□□ □□□□□.
- D. □□ □□□□ □□□□□.

E. master □□□□□□□ □□□□ □□□□.

F. Always Encrypted□ □□□□□.

Answer: B,E (LEAVE A REPLY)

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

Database-level firewall rules only apply to individual databases.

Reference:

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

NEW QUESTION: 3

DB1□□□ Azure SQL □□□□□□□ □□, □ □□□□□□□□ Sales□□ □□□ □□ □□□□ □□□□. Sales□ □□ □□ □□□ □□□□□ □□□□. □ □□ □□ □□□ □□□□ □□□. □□ Transact-SQL □ □ □□□ □□□?

```
ALTER CERTIFICATE Sales
```

A. WITH PRIVATE KEY (ENCRYPTION BY PASSWORD = '6YY9YcD!pV');

B. ALTER CERTIFICATE Sales WITH PRIVATE KEY (FILE = 'D:\importkeys\SalesNew, DECRYPTION BY PASSWORD = 'Mb^6BK&*w%');

```
ALTER CERTIFICATE Sales
```

C. WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'Mb^6BK&*w%', ENCRYPTION BY PASSWORD = '6YY9YcD!pV');

D. ALTER CERTIFICATE Sales WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'EWYx9Xk+\$#');

Answer: B (LEAVE A REPLY)

NEW QUESTION: 4

□□□□ □□ □□□□ □□ □□□□, □□ □□, □□ □□, □□□ □□ □ □□□ □□ □□□ □□ □□□ □□ □□□□ □□ □□□□. □□□□□ □□□ □□ □□ □□□□ □□□□.

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

□□□□ □□ □□ □□□ □□□□□ □□ □□□□ □□□□ □□□. □□ □□□□ □□ □□□ □□□□ □ □□ □□□□ □□□□□. □ □□ □□ □□ □□□□ □□□? □□□ □□□□□ □□□□□ □□□ □□□ □□□□□□□.

□□: □□ □□□ 1□□□□.

Microsoft

EventCategory: ▼
DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping: ▼
DimChannel
DimDate
DimEvent
FactEvents

TotalEvents: ▼
DimChannel
DimDate
DimEvent
FactEvents

Answer:

Microsoft

EventCategory: ▼
DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping: ▼
DimChannel
DimDate
DimEvent
FactEvents

TotalEvents: ▼
DimChannel
DimDate
DimEvent
FactEvents

Explanation:

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

Reference:

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

NEW QUESTION: 6

Which Azure service can be used to monitor and analyze the security and audit data for all the PaaS databases?

- A. Azure Key Vault
- B. Azure Event Hubs
- C. Azure Log Analytics

Answer: (SHOW ANSWER)

Explanation:

Scenario: Use a single dashboard to review security and audit data for all the PaaS databases.

With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics.

Note: Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/visualize/tutorial-logs-dashboards>

NEW QUESTION: 7

Azure Synapse Analytics uses Transact-SQL to query data in a data warehouse.

Which of the following is a valid Transact-SQL statement to create a table?

CREATE TABLE mytable (id INT, name VARCHAR(50), email VARCHAR(100), phone VARCHAR(20), address VARCHAR(255));

Transact-SQL uses the following syntax to create a table: CREATE TABLE table_name (column_name data_type, column_name data_type, column_name data_type, column_name data_type, column_name data_type);

CREATE TABLE mytable (id INT, name VARCHAR(50), email VARCHAR(100), phone VARCHAR(20), address VARCHAR(255));

Answer: C

Values

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

Answer Area

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

Answer:

Values

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

Answer Area

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

Explanation:

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

Box 1: HASH

Box 2: OrderDateKey

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

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management.

For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool>

NEW QUESTION: 8

Which of the following is a watermark delay metric? (Select two.)
A. The watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.
B. The watermark delay metric can rise due to:
1. Not enough processing resources in Stream Analytics to handle the volume of input events.

- A. The watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.
- B. The watermark delay metric can rise due to:
1. Not enough processing resources in Stream Analytics to handle the volume of input events.
- C. The watermark delay metric can rise due to:
1. Not enough processing resources in Stream Analytics to handle the volume of input events.
- D. The watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.

Answer: D (LEAVE A REPLY)

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: 9

□□ □□□ □□□ □□ □□ □□□□ □□□□ □□□. □ □□□□ □□ □□ □□□ □□□ □□ □□□ □□□□ □□□. □□□□ □□□ □□□□ □□□? □□□□□ □□□□□ □□□ □□□ □□□□□. □□: □□ □□□ 1□□□□.

Table type to store retail store data:

Table type to store promotional data:

Answer:

Table type to store retail store data:

Table type to store promotional data:

Explanation:

Table type to store retail store data:

Table type to store promotional data:

Box 1: Hash

Scenario:

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

A hash distributed table can deliver the highest query performance for joins and aggregations on large tables.


```
"resources": [  
  {  
    "type": "Microsoft.Sql/servers",  
    ...  
  },  
  {  
    "type": "Microsoft.Sql/servers/firewallRules",  
    "apiVersion": "2021-02-01-preview",  
    "name": "[concat(parameters('servers_a400n10102_name'), '/AllowAllWindowsAzureIps')]",  
    "dependsOn": [  
      "[resourceId('Microsoft.Sql/servers', parameters('servers_a400n10102_name'))]"  
    ],  
    "properties": {  
      "startIpAddress": "0.0.0.0",  
      "endIpAddress": "0.0.0.0"  
    }  
  }  
]
```

Explanation:



NEW QUESTION: 15

Azure SQL Database 100 100 1000000 1000.
100 10 100 10000.
100 1132, 10 16, 10 1, 10 1
100 10 10 100 100 1000000. 100 10 10000 10 1000 (76800)MB 1000 10 10000.
10 100 10000 1000. 10 1000 10000 1000 1000000 1000.
1000 10 100 100 1000 1000000? 10 1000 1000 10000 1000000.
100: 10 1000 100000.
A. 1000000 1000 100000000 1000000.
B. 10 10 10 10 1000 10000.

- C. □□□ □□□ □□□□□□.
- D. □□□□□□□□ □□□□ □□□□□□.
- E. □□ □□□□□□□ □□□□□□.

Answer: A,B,E ([LEAVE A REPLY](#))

NEW QUESTION: 16

Azure □□□□ SQLVM1□□□□ Azure □□ □□□ □□□□ SQL Server □□□□□ User1□□□□ □□□□ □□□□. SQLVM1□□ DB1□□□□ □□□□□□□□ □□□□□ □□□□. □□□(User!)□ DB1□ □□ □□□ □□□□ □□ □□□ □□□ □ □□□ □□ □□□. □ □□□□ □□ □□ □□□ □□□□ □□□□. User1□□ □□ □□ □□□□□□ □□□ □□□□ □□□□?

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

Answer: A ([LEAVE A REPLY](#))

DP-300-KR □□ □□□ □□□□□ □□ DumpTop □□ □□□□ □□□ DP-300-KR □□! DumpTop □ □□ **DP-300-KR** □□ □□□ □□□□□□□, DumpTop DP-300-KR □□ □□□ □□□□□□□□ □□□ □□ □□□□□. □□□□ □□□ □□□□ □□ DumpTop DP-300-KR □□□ □□□□□. <https://www.dumptop.com/Microsoft/DP-300-KR-dump.html> (410 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 17

SQL1□□□□ □□□ □ Azure SQL □□□ □□□□□ □□□□ □□□□□□ □□ □□ □□□□ □□□□□□. SQL1□□ □□□□ SQL Server □□□□□ □□□ □□ □□ □ □□□ □□ □□□□ □□ □□□. □ □□ □□□□ □□□□ □□ □ □□ □□□ □□□□□? 7. □□□□□ □□ □□□□ □□□ □□□ □□ □□□ □□□ □□□□□□.

Actions

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

Answer Area

>
<



Answer:

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

Explanation:

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

NEW QUESTION: 18

Microsoft SQL Server VM1 on Azure. You need to attach an ultra disk to VM1.

What should you do first? (Select two.)

Options: A. Stop and deallocate VM1. B. Attach the ultra disk. C. Set Enable Ultra disk compatibility to Yes. D. Start VM1. E. Resize VM1.

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

Answer:

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.

Explanation:

Answer Area
1 Stop and deallocate VM1.
2 Attach the ultra disk.
3 Set Enable Ultra disk compatibility to Yes .
4 Resize VM1.
5 Start VM1.

NEW QUESTION: 19

Azure SQL Columnstore indexes. Columnstore indexes are used to store and query data in a columnar format. Columnstore indexes are used to store and query data in a columnar format. Columnstore indexes are used to store and query data in a columnar format. Columnstore indexes are used to store and query data in a columnar format. Columnstore indexes are used to store and query data in a columnar format.

- A. Columnstore indexes are used to store and query data in a columnar format.
- B. Columnstore indexes are used to store and query data in a columnar format.
- C. Columnstore indexes are used to store and query data in a columnar format.
- D. Columnstore indexes are used to store and query data in a columnar format.
- E. Columnstore indexes are used to store and query data in a columnar format.

Answer: B,E (LEAVE A REPLY)

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

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



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: 21

Azure Synapse Analytics SQL (Pool1) Azure Data Lake Storage Gen2 (Account1) .
 Account1 .
 Pool1 .
 Transact-SQL .
 : .

```
CREATE EXTERNAL DATA SOURCE source1
WITH
( LOCATION = 'https://account1. .core.windows.net'
```

- blob
- dfs
- table

```
PUSHDOWN = ON
TYPE = BLOB_STORAGE
TYPE = HADOOP
```

Answer:

* 2020 2 15

* 2020 2 22

* 2020 2 29

* 2020 3 7

□□ □□ □□ □□ □□□□ □ □□□ □□□□ □□ □□□□ □□□□ □□□□□.

□□: □□ □□□ 1□□□□.

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

6 weeks
12 months
10 years

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

6 weeks
12 months
10 years

Answer:

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

6 weeks
12 months
10 years

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

6 weeks
12 months
10 years

Explanation:

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

	▼
6 weeks	
12 months	
10 years	

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



	▼
6 weeks	
12 months	
10 years	

NEW QUESTION: 26

4 Azure Automation Runbooks are configured to run on a daily basis. The Runbooks are configured to run on a daily basis. The Runbooks are configured to run on a daily basis. The Runbooks are configured to run on a daily basis.

- A. Azure Automation Runbook
- B. SQL Server Agent
- C. Azure SQL Database
- D. Azure SQL Database

Answer: A (LEAVE A REPLY)

Reference:

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

NEW QUESTION: 27

A Microsoft SQL Server 2019 database is configured with a 4TB database. The database is configured with a 4TB database. The database is configured with a 4TB database. The database is configured with a 4TB database.

- A. 100 GB
- B. 100 MB
- C. 100 KB
- D. Azure SQL Database

Answer: (SHOW ANSWER)

NEW QUESTION: 28

DB3 is a long-term retention backup. The backup is configured with a long-term retention backup. The backup is configured with a long-term retention backup. The backup is configured with a long-term retention backup.

- A. 100 GB
- B. 100 MB(LTR)

C. Point-in-Time Recovery (PITR)

D. Backup

Answer: (SHOW ANSWER)

NEW QUESTION: 29

500 users are using SQL1 in a Microsoft SQL Server environment.

The environment is on Azure SQL Database. The environment is configured for Always On Availability Groups.

The environment is configured for Always On Availability Groups. The environment is configured for Always On Availability Groups.

What is the maximum number of secondary replicas that can be configured in this environment?

A. 1

B. 2

C. 3

D. Always On Availability Groups

Answer: B (LEAVE A REPLY)

NEW QUESTION: 30

DB1 is a database on Azure SQL Database. The environment is configured for Always On Availability Groups.

The environment is configured for Always On Availability Groups. The environment is configured for Always On Availability Groups.

The environment is configured for Always On Availability Groups. The environment is configured for Always On Availability Groups.

What is the maximum number of secondary replicas that can be configured in this environment?

Options: 1, 2, 3, 4

A. sys.dm_exec_query_plan_stats

B. DB1 Lightweight_Query_Profiling

C. DB1 Last_Query_Plan_Stats

D. DB1 Lightweight_Query_Profiling

E. DB1 PARAMETER_SNIFFING

Answer: (SHOW ANSWER)

Last_Query_Plan_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm_exec_query_plan_stats.

Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT_QUERY_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>

NEW QUESTION: 31

Azure SQL Database 2017 Azure SQL Database environment is configured for Always On Availability Groups.

The environment is configured for Always On Availability Groups. The environment is configured for Always On Availability Groups.

Transact-SQL environment is configured for Always On Availability Groups. The environment is configured for Always On Availability Groups.

What is the maximum number of secondary replicas that can be configured in this environment?

A. Azure SQL Database

B. SQL Server Always On Availability Groups

C. Azure SQL Database

Answer Area

```
SELECT database_name = SQL1(),  
AVG(avg_memory_usage_percent) AS 'Average memory usage'  
MAX( [ ] ) AS 'Maximum CPU usage'  
FROM sys. [ ]  
Microsoft
```

Answer:
Answer Area

```
SELECT database_name = SQL1(),  
AVG(avg_memory_usage_percent) AS 'Average memory usage'  
MAX( [ ] ) AS 'Maximum CPU usage'  
FROM sys. [ ]  
Microsoft
```

Explanation:

```
Answer Area  
Microsoft  
SELECT database_name = SQL1(),  
AVG(avg_memory_usage_percent) AS 'Average memory usage'  
MAX( max_worker_percent ) AS 'Maximum CPU usage'  
FROM sys. dm_db_resource_stats
```


Statements	Yes	No
DB1 can be restored to a specific point in time 30 days ago.	<input type="radio"/>	<input checked="" type="radio"/>
DB1 can be restored from a weekly backup performed six months ago.	<input checked="" type="radio"/>	<input type="radio"/>
DB1 can be restored from a yearly backup performed six years ago.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Statements	Yes	No
DB1 can be restored to a specific point in time 30 days ago.	<input type="radio"/>	<input checked="" type="radio"/>
DB1 can be restored from a weekly backup performed six months ago.	<input checked="" type="radio"/>	<input type="radio"/>
DB1 can be restored from a yearly backup performed six years ago.	<input type="radio"/>	<input checked="" type="radio"/>

Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackupshorttermretentionpolicy?view=azps-7.2.0>
<https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackuplongtermretentionpolicy?view=azps-7.2.0>

NEW QUESTION: 39

50 Azure SQL databases are configured with the default backup retention policy. An administrator wants to ensure that all databases have a backup retention policy of 30 days. Which PowerShell cmdlet should be used to accomplish this task?

- A. Set-AzSqlDatabaseBackupShortTermRetentionPolicy
- B. Set-AzSqlDatabaseBackupLongTermRetentionPolicy
- C. Set-AzSqlDatabaseBackupRetentionPolicy
- D. Set-AzSqlDatabaseBackupPolicy

Answer: D (LEAVE A REPLY)

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: 40

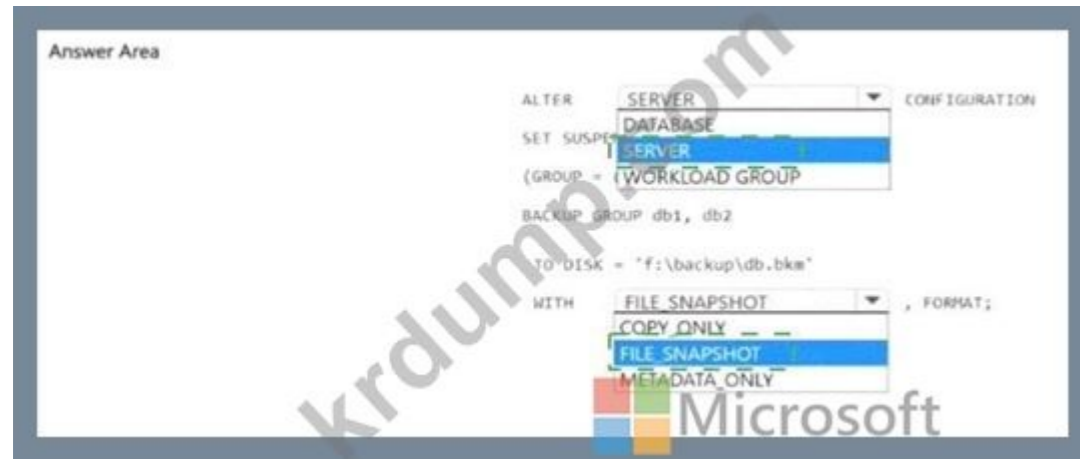
Azure SQL VM1 is a SQL Server instance with two databases, DB1 and DB2. DB1 is configured with a backup retention policy of 30 days. DB2 is configured with a backup retention policy of 90 days. Which PowerShell cmdlet should be used to ensure that both databases have a backup retention policy of 30 days?

I SQL 2019 2019 2019? 2019 20 2019 2019 2019 2019.
20: 20 2019 1019.

Answer Area



Answer:



Explanation:

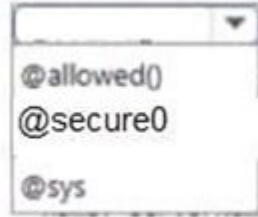


NEW QUESTION: 41

Azure 2019 2019 2019.
Bicep 2019 2019 Azure SQL Database 20 2019 20192019 2019.
20 2019 2019 2019 2019? 2019 20192019 20 20 2019 2019 2019 2019. (20: 20 2019 1019.)

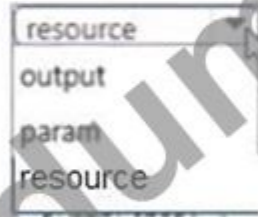
Answer Area

```
param serverName string = uniqueString('sql', resourceGroup().id)
param location string = resourceGroup().location
param administratorLogin string
```



A dropdown menu with a downward arrow on the right. The menu is open, showing three options: @allowed(), @secure0, and @sys.

```
param administratorLoginPassword string
```



A dropdown menu with a downward arrow on the right. The menu is open, showing four options: resource, output, param, and resource.

```
sqlServer 'Microsoft.Sql/servers/2024-05-01-preview' = {
```



```
administratorLogin: administratorLogin
administratorLoginPassword: administratorLoginPassword
}
}
```

Answer:

Answer Area



```
param serverName string = uniqueString('SQL', resourceGroup().id)
param location string = resourceGroup().location
param administratorLogin string

@allowed()
@secure0
@sys

param administratorLoginPassword string

resource sqlServer 'Microsoft.Sql/servers@2024-05-01-preview' = {
  output
  param
  resource

  administratorLogin: administratorLogin
  administratorLoginPassword: administratorLoginPassword
}
}
```


Explanation:

Answer Area

```

param serverName string = uniqueString('sql', resourceGroup().id)
param location string = resourceGroup().location
param administratorLogin string
@secure()
param administratorLoginPassword string
resource sqlServer 'Microsoft.Sql/servers@2024-05-01-preview' = {
  name: serverName
  location: location
  properties: {
    administratorLogin: administratorLogin
    administratorLoginPassword: administratorLoginPassword
  }
}

```




NEW QUESTION: 42

Vault1 is an Azure Key Vault. SQL1 is an Azure Virtual Machines SQL Server. SQL Server 1 has a database named DB1. Vault1 is configured to use D61 as a credential for SQL1 (TDE). SQL Server 1 is configured to use the credential. What should you do to ensure that the data in DB1 is encrypted at rest?

- ACTIONS
- On SQL1, create a symmetric key.
 - On SQL1, create a login from the symmetric key.
 - Create an Azure AD service principal and grant the service principal permissions for Vault1.
 - On SQL1, create an asymmetric key.
 - On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.
 - On SQL1, create a login from the asymmetric key.

Answer Area



Answer:




Actions

- On SQL1, create a symmetric key.
- On SQL1, create a login from the symmetric key.
- Create an Azure AD service principal and grant the service principal permissions for Vault1.
- On SQL1, create an asymmetric key.
- On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.
- On SQL1, create a login from the asymmetric key.


Answer Area

- Create an Azure AD service principal and grant the service principal permissions for Vault1.
- On SQL1, create an asymmetric key.
- On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.
- On SQL1, create a login from the asymmetric key.



Explanation:

On SQL1, create a symmetric key.	1 Create an Azure AD service principal and grant the service principal permissions for Vault1.
On SQL1, create a login from the symmetric key.	2 On SQL1, create an asymmetric key.
	3 On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.
	4 On SQL1, create a login from the asymmetric key.



NEW QUESTION: 43

Azure Databricks □□□□ □□□□.
 Databricks □□□□ □□ □□□ □□ □□□ □□□ □□□□ □□□□ □□□.
 □□ Databricks □□□□ □□□□ □□□□?

- A. □□□□
- B. □□
- C. DBFS
- D. SSH
- E. □□ □□

Answer: E (LEAVE A REPLY)

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:

NEW QUESTION: 44

Server1 is an Azure SQL Database server. You need to configure Server1 to audit data access to the database. The solution must meet the following requirements:

- The audit must be configured to capture all data access to the database.
- The audit must be configured to capture all data access to the database.
- The audit must be configured to capture all data access to the database.

Actions	Answer Area
<input type="checkbox"/> On Server1, enable diagnostics.	
<input type="checkbox"/> On Server1, run the create server audit specification T-SQL statement.	
<input type="checkbox"/> From the Azure portal, create an Azure Storage account.	
<input type="checkbox"/> From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.	
<input type="checkbox"/> On Server1, create a credential.	
<input type="checkbox"/> On Server1, run the create server audit T-SQL statement.	

Answer:

Actions	Answer Area
<input type="checkbox"/> On Server1, enable diagnostics.	<input type="checkbox"/> From the Azure portal, create an Azure Storage account.
<input type="checkbox"/> On Server1, run the create server audit specification T-SQL statement.	<input type="checkbox"/> From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
<input type="checkbox"/> From the Azure portal, create an Azure Storage account.	<input type="checkbox"/> On Server1, create a credential.
<input type="checkbox"/> From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.	<input type="checkbox"/> On Server1, run the create server audit T-SQL statement.
<input type="checkbox"/> On Server1, create a credential.	
<input type="checkbox"/> On Server1, run the create server audit T-SQL statement.	

Explanation:

Actions

- On Server1, enable diagnostics.
- On Server1, run the create server audit specification T-SQL statement.

Answer Area

- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, run the create server audit T-SQL statement.

NEW QUESTION: 45

SQL1 is an Azure SQL Database instance. You need to configure SQL1 to send diagnostic data to an Azure Storage account. Which two actions should you perform? (Select two.)

- A. Create a storage account and create a storage container.
- B. Create a storage account and create a storage endpoint.
- C. Microsoft SQL Server Enterprise Edition is required.
- D. Create a storage account and create a storage endpoint.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 46

10 SQLMI instances are running on an Azure SQL Database instance. You need to configure the instances to send diagnostic data to Log Analytics. Which two actions should you perform? (Select two.)

- Create a storage account and create a storage endpoint.
- Create a storage account and create a storage endpoint.

- A. Create a storage account and create a storage endpoint.
- B. Create a storage account and create a storage endpoint.
- C. SQLMI instances are required.
- D. SQLMI instances are required.

Answer: A (LEAVE A REPLY)

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure?tabs=azure-portal#configure-the-streaming-export-of-diagnostic-telemetry>

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')))]"
          ]
        ]
      ]
    }
  ]
}
```

Answer:

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')))]"
              ]
            ]
          ]
        }
      ]
    }
  ]
}
```

Explanation:

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

```

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

NEW QUESTION: 49

Scenario: You have a Microsoft Azure subscription named Contoso. The subscription contains two servers named Server1 and Server2. Server1 is running Microsoft SQL Server 2017. Server2 is running Microsoft SQL Server 2019. You have a Microsoft Azure SQL Database named Database1. Database1 is running on Server1. You need to restore Database1 to Server2. The restore must use the RESTORE Transact-SQL command and the REPLACE option. What should you do? (Select two.)

- A. Restore Database1 from Server1 to Server2.
- B. Restore Database1 from Server1 to Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Answer: B (LEAVE A REPLY)

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

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

ence:



NEW QUESTION: 51

Q: Which Azure SQL service tier provides the highest availability? A: Premium. 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>

- A. Premium
- B. Business Critical

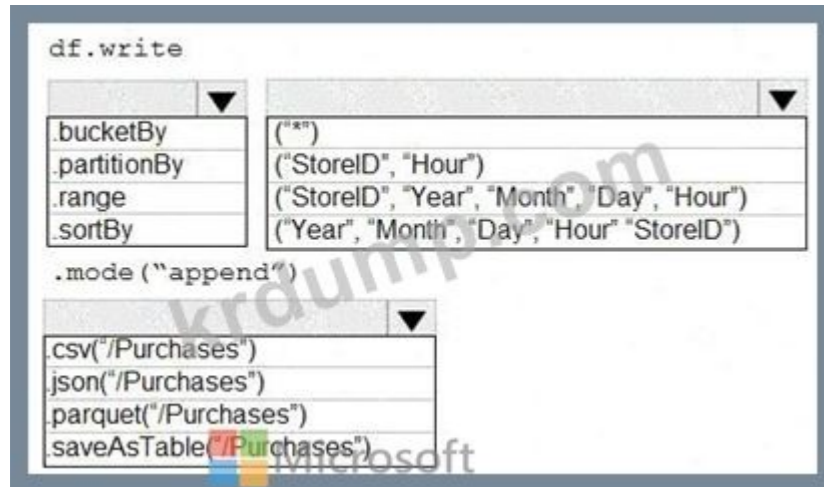
Answer: A (LEAVE A REPLY)

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>

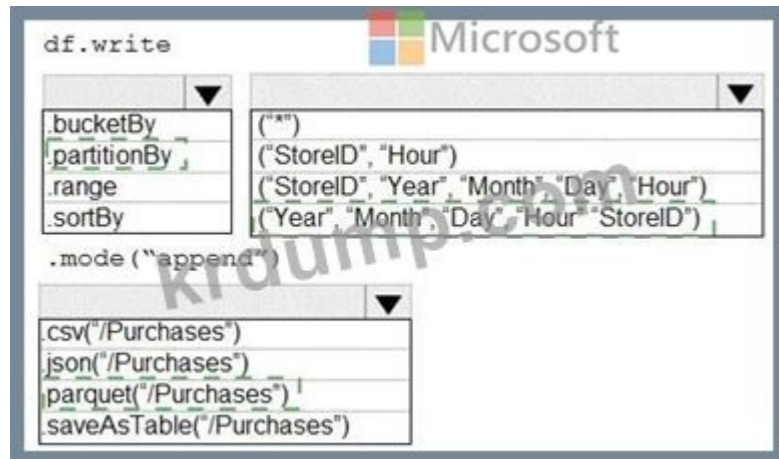
NEW QUESTION: 52

Q: Which Azure Databricks configuration property is used to specify the storage location for the cluster logs? A: StoreLogsIn. StoreLogsIn is a boolean property that, when set to true, stores cluster logs in the default storage location for the cluster. The default storage location is the Azure Storage account specified in the cluster configuration. Reference: <https://docs.microsoft.com/en-us/azure/databricks/cluster-config-properties>

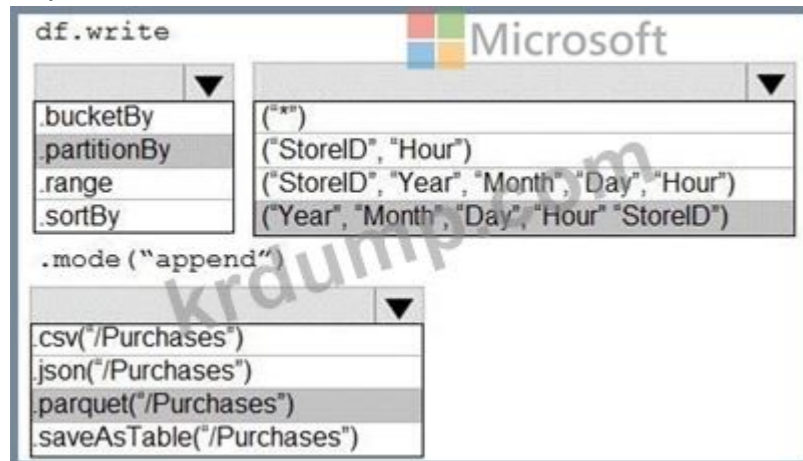
□□□ □□□ □□□□ □□□? □□□ □□□□□ □□□□□ □□□ □□□ □□□□.
 □□: □□ □□□ 1□□□□.



Answer:



Explanation:



Box 1: .partitionBy

Example:

```
df.write.partitionBy( " y " , " m " , " d " )
```

```
mode(SaveMode.Append)
```

```
parquet( " /d ata/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-partitions-with-no-new-data>

NEW QUESTION: 53

DB1 is a Microsoft SQL Server database on Azure. DB1 contains CustomerPII data. An auditing policy is defined for the database. Which of the following is true?

- A. The auditing policy applies to all existing and newly created databases on the server.
- B. The auditing policy applies to the database, regardless of the database auditing settings.
- C. 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.
- D. The auditing policy applies to the database, but not to the server.

Answer: A,C (LEAVE A REPLY)

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: 54

Azure SQL Database uses Windows Event Tracing for Windows (ETW) to collect performance and diagnostic data. Which of the following is true?

- A. The auditing policy applies to all existing and newly created databases on the server.
- B. The auditing policy applies to the database, regardless of the database auditing settings.
- C. 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.
- D. The auditing policy applies to the database, but not to the server.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 55

DB1 is a Microsoft SQL Server 2016 database on Azure. DB1 contains CustomerPII data. An auditing policy is defined for the database. Which of the following is true?

- A. Azure SQL Database uses Windows Event Tracing for Windows (ETW) to collect performance and diagnostic data.
- B. Azure SQL Database uses Windows Event Tracing for Windows (ETW) to collect performance and diagnostic data.

C. PostgreSQL Azure

D. Azure SQL Server

Answer: B (LEAVE A REPLY)

Azure SQL Database single database supports Data Sync.

Reference:

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

NEW QUESTION: 56

Pool1 Azure Synapse Analytics SQL DB1

DB1 Table

Table1

Synapse Studio

A. Pool1 sys.dm_pdw_nodes_db_partition_stats

B. DBCC CHECKALLOC

C. Pool1 DBCC CHECKALLOC

D. sys.dm_pdw_nodes_db_partition_stats

Answer: D (LEAVE A REPLY)

Use sys.dm_pd w_nodes_db_partition_stats to analyze any skewness in the data.

Reference:

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

NEW QUESTION: 57

Azure Data Lake Storage Gen2 Azure Data Factory

LastUpdatedDate

4

-

-

A.

B.

C.

D.

Answer: A (LEAVE A REPLY)

The Tumbling window trigger supports backfill scenarios. Pipeline runs can be scheduled for windows in the past.

:

https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers

NEW QUESTION: 58

SERVER1

□□□□□□ □□□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□□.
□□: □□ □□□ 1□□□□□.

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

Answer:

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

Explanation:

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

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.

Premimum 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: 59

Azure 00 000 00 0 00 SQL Server 00000 000 00 000 000 0000.

00 00000 00 00 00000 000 Azure Marketplace SQL Server 2019 Enterprise 0000 0000 00000000. 000000 FCI10000 000 00 00 0000 00000(FCI)0 000 00000 0 000.

000000 00000000 FCI10 0000 0 0000 00 0000. 00000 00 00 0000 00000 0000.

* 0000 SLA0 000000

0000 00000000.

0000 00000 0000?

A. 00 00000 00(VNN) 0000

B. 00 Azure 00 0000

C. 00 00000 00(DNN) 0000

D. Azure 00 00 0000

Answer: C (LEAVE A REPLY)

NEW QUESTION: 60

Azure SQL Database0 00 00(Server1)0 00, 00 00000 Pool0000 0000 00 00000. 00 Watcher 10000 00000000 00000 00000. 00 KQL 0000 000000.

```
.set logicalServer = @"Server1";
.set elasticPoolName = @"Pool1";
.set replicaType = "Primary";
.set duration = 20h;
sqldb_elastic_pool_resource_utilization
where sample_time_utc > ago(duration)
where logical_server_name = ~logicalServer
where replica_type = ~replicaType
project sample_time_utc,
    avg_cpu_percent,
    avg_instance_cpu_percent,
    avg_data_io_percent,
    avg_log_write_percent,
    max_worker_percent
sort by sample_time_utc desc
header timespan:
```

00 0 0000 00000 '0'0 00000, 0000 0000 '0000'0 00000000.

00: 00 0000 100000.

ACTIONS

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.

ANSWER AREA



Answer:

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

Explanation:

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"/>

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: 64

Azure SQL .

.



.

: .

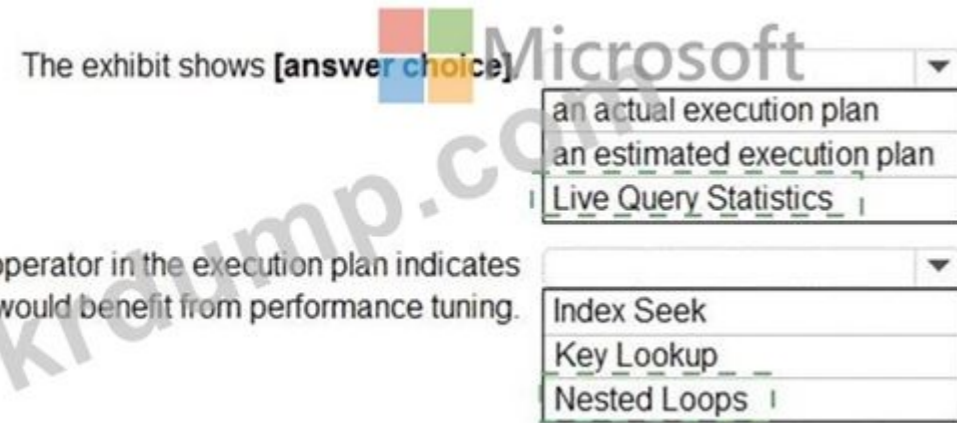
The exhibit shows

The operator in the execution plan indicates that the query would benefit from performance tuning.

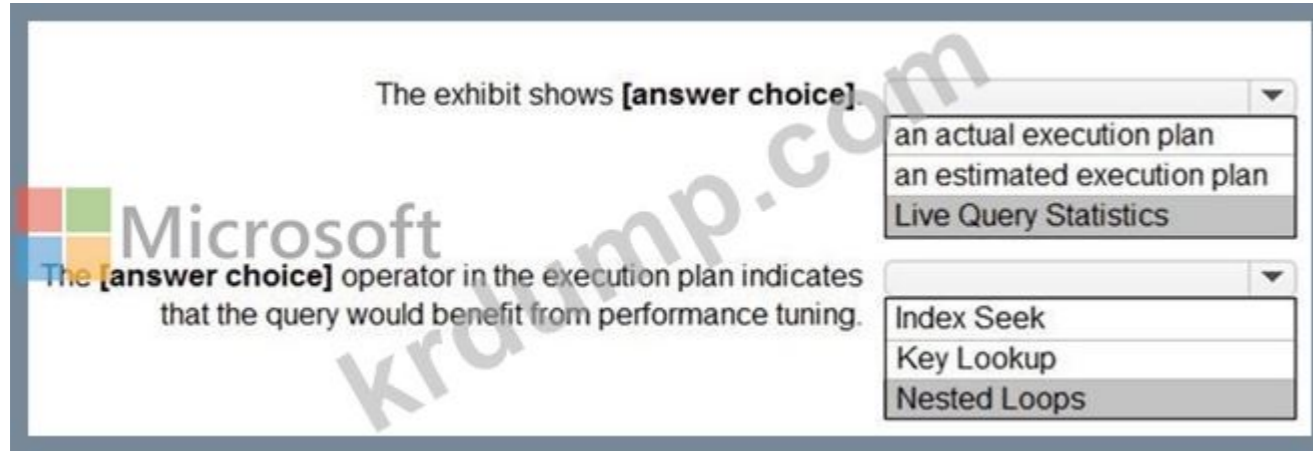
Answer:

The exhibit shows [answer choice]

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



Explanation:



Reference:

[https:// docs.microsoft.com/en-us/sql/relational-databases/performance/li ve-query-statistics?view=sql-server- ver15](https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15)

NEW QUESTION: 65

Azure Virtual Machines DB1 is a Microsoft SQL Server database. DB1 is an Azure SQL database. DB1 is a Microsoft SQL Server database. DB1 is an Azure SQL database. DB1 is a Microsoft SQL Server database. DB1 is an Azure SQL database.

- A. Microsoft SQL Server
- B. Azure SQL Database
- C. Microsoft SQL Server
- D. Azure SQL Database

Answer: B (LEAVE A REPLY)

Reference:

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

NEW QUESTION: 66

DB1 is a Microsoft SQL Server database. DB1 is an Azure SQL database. DB1 is a Microsoft SQL Server database. DB1 is an Azure SQL database. DB1 is a Microsoft SQL Server database. DB1 is an Azure SQL database.



Which of the following actions would you take to reduce the memory usage of the query plan?

- A. DBCC FREEPROCCACHE
- B. ALTER INDEX [index_name] REORGANIZE
- C. ALTER INDEX [index_name] REBUILD
- D. ALTER INDEX [index_name] SET (IGNORE_DUP_KEY = OFF)

Answer: (SHOW ANSWER)

Reference:

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

NEW QUESTION: 67

A user named user1 is running a query on an Azure SQL Database. The query is running on a server pool named DB1. The query is running on a server named LTR. The query is running on a server named LTR. The query is running on a server named LTR. User1 is running the query on a server named LTR. Which of the following actions would you take to reduce the memory usage of the query plan?

- A. SQL Server
- B. SQL
- C.
- D.

Answer: (SHOW ANSWER)

NEW QUESTION: 68

Azure SQL Database is a fully managed relational database service. It is a cloud-based service that allows you to create and manage a SQL database in the cloud. It is a fully managed service, which means that Microsoft handles all the infrastructure and maintenance tasks. You can create a database in the cloud and manage it using the same tools and techniques as you would use to manage a database on-premises. Azure SQL Database is a fully managed service that allows you to create and manage a SQL database in the cloud. It is a fully managed service, which means that Microsoft handles all the infrastructure and maintenance tasks. You can create a database in the cloud and manage it using the same tools and techniques as you would use to manage a database on-premises.

USE master;

ALTER DATABASE [DB1] SET

▼
OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY

WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

▼
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST

WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

▼
MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;

GO



Answer:

USE master;

ALTER DATABASE [DB1] SET

▼
OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY

WITH ROLLBACK IMMEDIATE

GO

DBCC CHECKDB ('DB1',

▼
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST

WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

▼
MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;

GO



Explanation:

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

dbo.FactInternetSales □□□□ □□□□ □□□□ □□□ □□□□□?

- A. □□□□ 10,000□ □□□ □□ □□□□ □□□□.
- B. □□ □□□□ □□□□ □□□□□.
- C. □□□□ □□□ □□ □□□ □□□□□.
- D. □□□□ □□□□□□□.

Answer: D (LEAVE A REPLY)

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

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

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

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

Reference:

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

<https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage- distributed-data-skew.md>

NEW QUESTION: 70

sqldb1□□□ Azure SQL □□□□□□□□ □□□□.

sqldb1□ □□□ □ □□ □□□ □□□□□ □□□.

□□□ □□□□□□□?

- A. DBCC □□□□□□□□
- B. sp_clean_db_free_space
- C. sp_clean_db_file_free_space
- D. DBCC □□ □□


```

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

```

...
 ☐☐ ☐ ☐☐☐ ☐☐, ☐☐☐ ☐☐☐☐ '☐'☐ ☐☐☐☐, ☐☐☐ ☐☐☐ '☐☐☐'☐ ☐☐☐☐☐☐.
 ☐☐: ☐☐ ☐☐☐ 1☐☐☐☐.

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

Answer:

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

Explanation:

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

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/purchasing-models>

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

NEW QUESTION: 74

VM1 and VM2 are virtual machines running Microsoft SQL Server. VM1 has a database named db1. VM2 has a database named db2. Both databases are in the same availability group (DAG). The DAG is configured with a primary replica on VM1 and a secondary replica on VM2. The secondary replica on VM2 is in a state of 'SUSPECTED'. You need to restore the secondary replica on VM2. What should you do first?

Actions

- From the master database, run CREATE ENDPOINT.
- From db1, run CREATE CERTIFICATE.
- From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.
- From the master database, run CREATE CERTIFICATE.
- From db1, run CREATE ENDPOINT.

Answer Area

Answer:



Actions

- From the master database, run CREATE ENDPOINT.
- From db1, run CREATE CERTIFICATE.
- From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.
- From the master database, run CREATE CERTIFICATE.
- From db1, run CREATE ENDPOINT.

Answer Area

- From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.
- From the master database, run CREATE CERTIFICATE.
- From db1, run CREATE ENDPOINT.

Explanation:

Actions

- From the master database, run CREATE ENDPOINT.
- From db1, run CREATE CERTIFICATE.

Answer Area

- From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.
- From the master database, run CREATE CERTIFICATE.
- From db1, run CREATE ENDPOINT.

NEW QUESTION: 75

□□□□ □□□□ Always On □□□ □□□ □□□□ Azure □□ □□□ SQL Server □□□□ □ □□ □□□ □□□□□.

□□ □□ □□□ □□□□ □□ □□□□ □□□ □□□ □□□.

* □□ 99.95%□ □□□ □□ □□(SLA)□ □□□□□.

* □□ □□ □□ □□□□□□□ □□□□□ □□□□□.

* □□□□□□□ □□ □□ □□□ □□□□□□□

□□□ □□□□□□□□□?

A. □□ □□□ □□□ □□□ □□□□□. □ □□ □□□ □□□ □□□ □□□□□. □ □□ □□□ □□ □□ □□□ □□□□□.

B. □□ □□ □ □□ □□□ □□ □□□ □□□□□. □ □□ □□□ □□□ □□□ □□□□□. □ □□ □□□ □□ □□□ □□□ □□□□□.

C. □ □□ □□ □□□ □ □□ □□□ □□□ □□□□□. □ □□ □□□ □□□ □□□ □□□ □□□□□. □ □□ □□□ □□ □□□ □□□ □□□□□.

D. □□ □□□ □ □□ □□□ □□□ □□□□□. □ □□ □□□ □□□ □□□ □□□ □□□□□. □ □□ □□□ □□ □□ □□□ □□□□□.

Answer: A (LEAVE A REPLY)

To get VMs as close as possible, achieving the lowest possible latency, you should deploy them within a proximity placement group. <https://learn.microsoft.com/en-us/azure/virtual-machines/co-location>

NEW QUESTION: 76

□□□□□ Microsoft SQL Server 2022 □□□□□ DB1□□□□ □□□□□□□□ □□□□. Azure □□□□ SQLDB1□□□□ Azure SQL □□□□□□□□ □□□□. DB1□ SQLDB1□ □□□□ □□□□. □□ □□□ □□□ □

□□□ □□□?

A. □□

B. □□

C. □□□□□

Steps

- ⌵ Enable a system-assigned managed identity.
- ⌵ Enable a system-assigned service principal.
- ⌵ Grant admin consent to an app registration in Microsoft Entra.
- ⌵ Implement Microsoft Entra Cloud Sync.

Answer Area

Step 1: Enable Microsoft Entra authentication on Instance1.

Step 2: Implement Microsoft Entra Cloud Sync.

Step 3: Enable a system-assigned service principal.



NEW QUESTION: 78

Azure Data Lake Storage Gen 2 is a cloud storage service that supports various authorization mechanisms. Which of the following is NOT supported by Azure Data Lake Storage Gen 2?

- A. Shared Key
- B. Shared access signature (SAS)
- C. Role-based access control (Azure RBAC)
- D. Azure AD(Azure Active Directory) ID

Answer: (SHOW ANSWER)

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

- What resources the client may access.
- What permissions they have to those resources.
- How long the SAS is valid.

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

- * Shared Key authorization
- * Shared access signature (SAS) authorization
- * Role-based access control (Azure RBAC)
- * Access control lists (ACL) Data Lake Storage Gen2 supports the following authorization mechanisms:
- * Shared Key authorization
- * Shared access signature (SAS) authorization
- * Role-based access control (Azure RBAC)
- * Access control lists (ACL)

Reference:

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

NEW QUESTION: 79

Project1 is a storage account in Azure. Project2 is a storage account in Azure. Project1 is connected to Project2 via a storage sync service. Which of the following is NOT supported by Project1?

- * 25TB of storage capacity.
- * 25TB of storage capacity.
- * Project2 can be used as a sync target.
- * Project2 can be used as a sync source.

* 10TB of storage for SQL databases.

* 10TB of storage for SQL databases.

* 10TB of storage for SQL databases.

* 10TB of storage for SQL databases.

Which Azure SQL service tier is the most cost-effective for a workload that requires 10TB of storage and is not mission-critical?

Options: Azure SQL Database, Azure SQL Managed Instance in the Business Critical service tier, Azure SQL Managed Instance in the General Purpose service tier, SQL Server on Azure Virtual Machines.

Answer: Azure SQL Database.

Answer Area

Project1:

Project2:

Answer:

Answer Area

Project1:

Project2:

Explanation:

Answer Area

Project1: SQL Server on Azure Virtual Machines

Project2: Azure SQL Managed Instance in the General Purpose service tier



NEW QUESTION: 80

□□ □□□ □□ Azure □□□□ Azure SQL □□□□□□□ □□□□□□□ □□□□.

Microsoft Azure Search resources, services, and docs (G+)

liveid238@outlook.com DEFAULT DIRECTORY

Home > SQL databases > Create SQL Database >

Configure

Feedback

Compute Hardware

Click "Change configuration" to see details for all hardware generations available including memory optimized and compute optimized options

Hardware Configuration

Gen5
up to 40 vCores, up to 120 GB memory
[Change configuration](#)

Max vCores

1 2 4 6 8 10 12 14 16 18 20 24 32 40

Min vCores

0.75 1 1.25 1.5 1.75 2 2.25 2.5 3 4 5 6

2.25 GB MIN MEMORY 18 GB MAX MEMORY

Auto-pause delay

The database automatically pauses if it is inactive for the time period specified here, and automatically resumes when database activity recurs. Alternatively, auto-pausing can be disabled.

Enable auto-pause

Days: Hours: Minutes:

Data max size [Ⓢ]

1 GB 800 GB 1.5 TB

240 GB LOG SPACE ALLOCATED

Cost summary

Gen5 - General Purpose (GP_5, Gen5_0)

Cost per GB (in USD) **0.12**


Max storage selected (in GB) **x 1040**

ESTIMATED STORAGE COST / MONTH 119.60 USD

COMPUTE COST / VCORE / SECOND ¹ 0.000145 USD

NOTES

¹ Serverless databases are billed in vCores based on a combination of CPU and memory utilization. [Learn more about serverless billing](#)



□□ □□□ □□□ □□□ □ □□□ □□□□ □□ □□□□ □□□□ □□□□□.

□□: □□ □□□ 1□□□□.

After four hours of inactivity, the database requires [answer choice] to resume operations for new activities.

Microsoft
no extra time
up to 10 minutes
up to one minute

The database configuration reduces the cost of [answer choice] usage patterns.

Microsoft
intermittent and unpredictable
regular and high
steady and low

Answer:

Microsoft
After four hours of inactivity, the database requires [answer choice] to resume operations for new activities.
no extra time
up to 10 minutes
up to one minute
The database configuration reduces the cost of [answer choice] usage patterns.
intermittent and unpredictable
regular and high
steady and low

Explanation:

Microsoft
After four hours of inactivity, the database requires [answer choice] to resume operations for new activities.
no extra time
up to 10 minutes
up to one minute
The database configuration reduces the cost of [answer choice] usage patterns.
intermittent and unpredictable
regular and high
steady and low

Reference:

[https://docs.microsoft.com/en-us/azure/azure-sql /datab ase/serverless-tier-overview](https://docs.microsoft.com/en-us/azure/azure-sql/datab ase/serverless-tier-overview)

NEW QUESTION: 81

DB1 Azure SQL Table1

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

Table1 is Always Encrypted. How many columns are encrypted?
 How many columns are encrypted? How many columns are encrypted?
 How many columns are encrypted?

- A. 3
- B. 5
- C. 2
- D. 1
- E. 4

Answer: D,E (LEAVE A REPLY)

NEW QUESTION: 82

FileTable Filestream Microsoft SQL Server Azure SQL Database?
 Azure SQL Database?
 Azure SQL Database?


- A. Azure SQL Database
- B. Azure File Share SQL Server
- C. Azure SQL Database
- D. MySQL Azure Database


Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 83

Azure SQL Server? How many columns are encrypted?
 How many columns are encrypted?



Answer Area 

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

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

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

- Max Size (MB)
- Query Store Capture Mode
- Size Based Cleanup Mode
- Operation Mode (Requested)**

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

- all
- none of the
- a selective set of

Answer:

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

- Max Size (MB)
- Query Store Capture Mode
- Size Based Cleanup Mode
- Operation Mode (Requested)**

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

- all
- none of the
- a selective set of

Explanation:

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

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

NEW QUESTION: 84

server1 [redacted] db1 [redacted] Azure SQL [redacted] [redacted].
 Intelligent Insights [redacted] [redacted] [redacted] [redacted] [redacted] [redacted].
 [redacted] [redacted] [redacted] [redacted] [redacted] [redacted].
 [redacted] [redacted] [redacted] [redacted]?

- A. DBCC SQLPERF [redacted] [redacted].
- B. DBCC DBREINDEX [redacted] [redacted].
- C. db1 [redacted] [redacted] [redacted] [redacted].

D. db1 is not a serverless database.

Answer: C (LEAVE A REPLY)

Reference:

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

NEW QUESTION: 85

Azure SQL Server instances 500 are in Azure. An Azure SQL instance 500 is serverless. The instance is not a serverless instance. Which of the following is true?

* The instance is not a serverless instance.

* Azure Defender is not enabled for the instance.

* The instance is not a serverless instance.

The instance is not a serverless instance.

Which of the following is true? The instance is not a serverless instance.

Answer: D, E (LEAVE A REPLY)

A. Azure SQL Server

B. Azure Automation

C. Azure Automation

D. Azure Policy

E. Azure SQL Server

Answer: D, E (LEAVE A REPLY)

NEW QUESTION: 86

An Azure resource group contains the following resources:

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

App1 is connected to db1. The instance is not a serverless instance. Which of the following is true?

A. db1 is not a serverless instance.

B. The instance is not a serverless instance.

C. db1 is not a serverless instance.

D. db1 is not a serverless instance.

Answer: (SHOW ANSWER)

NEW QUESTION: 87

DB1 is an Azure SQL database instance.

Azure Portal is used to view the execution plan for a query.

Which of the following is true?

A. SET SHOWPLAN_ALL Transact-SQL is used to view the execution plan.

B. DB1 is not a serverless instance.

Objects

- a database encryption key in the master database
- a master key in DB1
- a certificate in DB1
- a master key in the master database
- a certificate in the master database
- a database encryption key in DB1

Answer Area

Answer:

Objects

- a database encryption key in the master database
- a master key in DB1
- a certificate in DB1
- a master key in the master database
- a certificate in the master database
- a database encryption key in DB1

Answer Area

- a master key in the master database
- a certificate in the master database
- a database encryption key in DB1

Explanation:

Objects

- a database encryption key in the master database
- a master key in DB1
- a certificate in DB1

Answer Area

- a master key in the master database
- a certificate in the master database
- a database encryption key in DB1

NEW QUESTION: 91

ResearchDB1 is a database. You need to create a user in ResearchDB1. Which T-SQL statement should you use?

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

Answer: (SHOW ANSWER)

Scenario: Authenticate database users by using Active Directory credentials.

- A. 10
- B. 1000000
- C. 100000 1000
- D. 10000

Answer: C (LEAVE A REPLY)

NEW QUESTION: 95

Which Microsoft SQL Server feature is supported by Azure SQL Database? (Select all that apply.)

- A. Azure Database for MySQL
- B. Azure VM-hosted SQL Server
- C. Azure SQL Database Elastic Pools
- D. Azure SQL Database Federated Query

Answer: B (LEAVE A REPLY)

NEW QUESTION: 96

Azure Synapse Analytics DW1000L resource class is used for a query. The query is a simple SELECT statement. The query is executed on a node pool with 100 nodes. The query is executed on a node pool with 100 nodes. The query is executed on a node pool with 100 nodes. The query is executed on a node pool with 100 nodes.

- A. The performance capacity of a query is determined by the user's resource class.
- B. DW1000L resource class has the highest performance capacity.
- C. Smaller resource classes reduce the maximum memory per query, but increase concurrency.
- D. Larger resource classes increase the maximum memory per query, but reduce concurrency.

Answer: C (LEAVE A REPLY)

The performance capacity of a query is determined by the user's resource class. Smaller resource classes reduce the maximum memory per query, but increase concurrency. Larger resource classes increase the maximum memory per query, but reduce concurrency. Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workloadmanagement>

NEW QUESTION: 97

Azure SQL Database is used for a query. The query is a simple SELECT statement. The query is executed on a node pool with 100 nodes. The query is executed on a node pool with 100 nodes. The query is executed on a node pool with 100 nodes. The query is executed on a node pool with 100 nodes.

Answer: C (LEAVE A REPLY)

Actions

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the CHANGETABLE() function.

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON
```

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the CHANGE_TRACKING_CURRENT_VERSION() function.

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_table
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_db
```

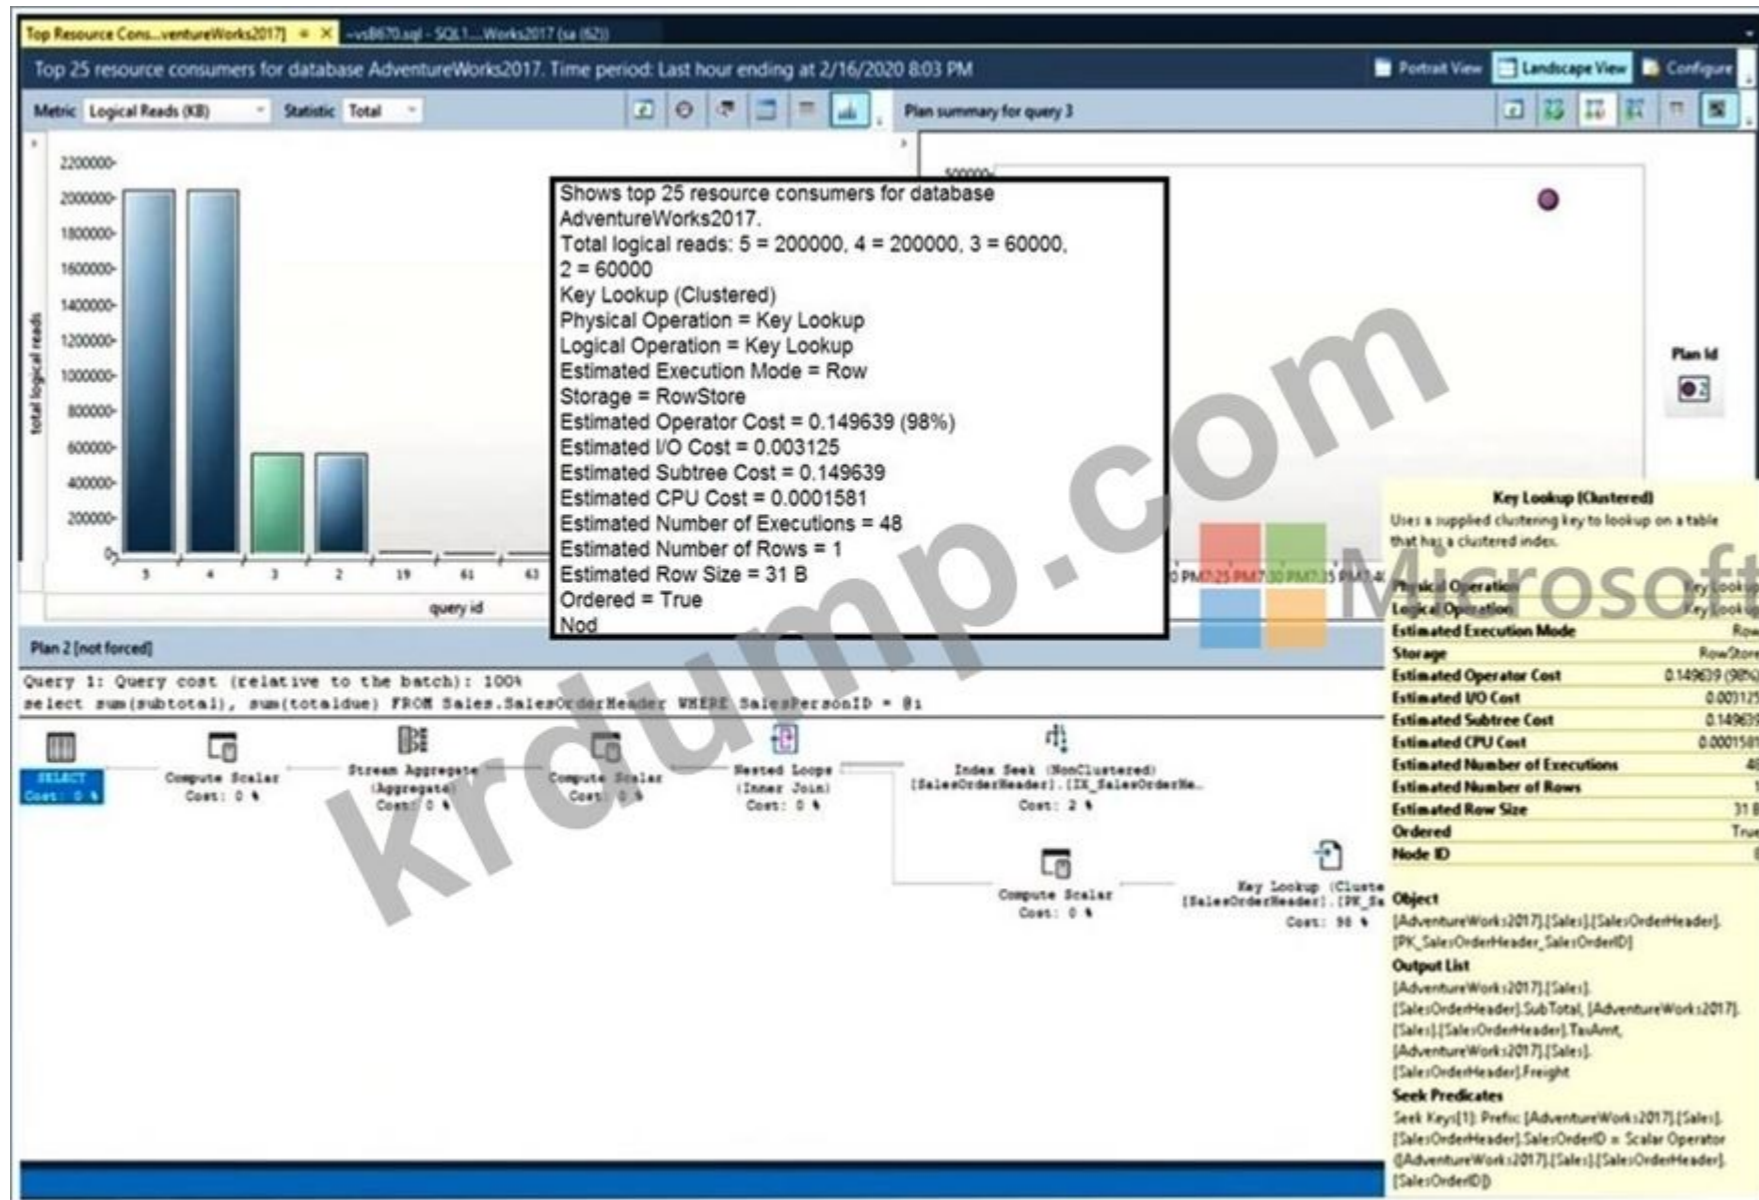
Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING
```

Answer Area

Answer:





□□ □ □□ □□, □□□ □□□□ '□'□ □□□□, □□□ □□□ '□□□'□ □□□□□□.
 □□: □□ □□□ 1□□□□.

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<input type="radio"/>	<input type="radio"/>
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	<input type="radio"/>	<input type="radio"/>
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<input type="radio"/>	<input type="radio"/>
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	<input type="radio"/>	<input checked="" type="radio"/>
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<input type="radio"/>	<input type="radio"/>
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	<input type="radio"/>	<input checked="" type="radio"/>
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	<input checked="" type="radio"/>	<input type="radio"/>

Reference:

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

NEW QUESTION: 99

Customers use dynamic data masking in Azure Synapse Analytics. Which of the following is a valid credit card masking method? (Select two.)

- A. XXXX-XXXX-XXXX-XXXX
- B. XXX-XXXX-XXXX-XXXX
- C. XXX-XXXX-XXXX-XXXX
- D. XXX-XXXX-XXXX-XXXX

Answer: B (LEAVE A REPLY)

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: 100

MI1 is an Azure SQL Managed Instance. Which of the following is a valid credit card masking method? (Select two.)

XXXX-XXXX-XXXX-XXXX?

- A. Azure Portal
- B. Transact-SQL `force_last_good_plan`
- C. REST API `AutomaticTuningServerMode`

Answer: B (LEAVE A REPLY)

NEW QUESTION: 101

Microsoft SQL Server 2022 Enterprise is installed on a server. The server is connected to the Internet. The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure.

The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure.

* The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure.

The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure.

The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure. The server is running on a virtual machine in Azure.

A. Azure SQL Database

B. Azure SQL

C. Azure SQL Server

D. Azure SQL

Answer: A (LEAVE A REPLY)

NEW QUESTION: 102

Microsoft SQL Server is installed on Server1 in Azure. Server1 is connected to the Internet. Server1 is connected to the Internet. Server1 is connected to the Internet.

ASVR1 is connected to the Internet. ASVR1 is connected to the Internet. ASVR1 is connected to the Internet.

SQL is connected to the Internet. SQL is connected to the Internet. SQL is connected to the Internet.

ASVR1 is connected to the Internet. ASVR1 is connected to the Internet. ASVR1 is connected to the Internet.

SQL is connected to the Internet. SQL is connected to the Internet. SQL is connected to the Internet.

ASVR1 is connected to the Internet. ASVR1 is connected to the Internet. ASVR1 is connected to the Internet.

Answer Area



Microsoft

Sync data by using:

Azure Private Link	▼
A service endpoint	
Azure Private Link	
The SQL Data Sync private link	

Set the database isolation level to:

SNAPSHOT ISOLATION	▼
REPEATABLE READ	
SERIALIZABLE	
SNAPSHOT ISOLATION	

Answer:

Quorum model:

▼
Cloud witness
Disk witness
File share witness

Azure resource for the availability group listener:

▼
Azure Application Gateway
Azure Basic Load Balancer

Answer:

Quorum model:	▼
	Cloud witness
	Disk witness
	File share witness
Azure resource for the availability group listener:	▼
	Azure Application Gateway
	Azure Basic Load Balancer

Explanation:

Quorum model:	▼
	Cloud witness
	Disk witness
	File share witness
Azure resource for the availability group listener:	▼
	Azure Application Gateway
	Azure Basic Load Balancer

Scenario: Business Requirements

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

Box 1: Cloud witness

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

Box 2: Azure Basic Load Balancer

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

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

reference:

<https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness>

https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/

NEW QUESTION: 105

1000 Azure Data Factory objects.

00, 00 00 0000 00 000 0000 0 000000 0000 0000 000. Data Factory00 0000 000 000 0 0000 000 0 0000 000 0 000 000.

0 000000 000 0000 000?

- A. 00
- B. 000 00
- C. 00 00 ID
- D. 000 00
- E. 00 ID

Answer: (SHOW ANSWER)

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.

Reference:

<https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/>

NEW QUESTION: 106

Azure SQL 0000000 0000.

00 00 000 0000 00 00 0000 0000 000.

* 000000 00 00 0 0000000 0000 00000 000 0000000.

* 000 000 0000000

0000 000 0000 000?

- A. 00 0000 00
- B. Azure 000 00
- C. 00 000 00
- D. Azure 00

Answer: A (LEAVE A REPLY)

DP-300-KR 00 000 000000 00 DumpTop 00 0000 000 DP-300-KR 00! DumpTop 0 00 **DP-300-KR** 00 000 000000, DumpTop DP-300-KR 00 000 00000000 000 00 00000. 0000 000 0000 00 DumpTop DP-300-KR 000 00000. <https://www.dumptop.com/Microsoft/DP-300-KR-dump.html> (410 Q&As Dumps, **30%OFF Special Discount: KrDump**)

NEW QUESTION: 107

PostgreSQL 0000000 0000000000 000 0000 000?

- A. Azure Data Box
- B. AzCopy
- C. Azure 000000 000000 000
- D. Azure 000 00

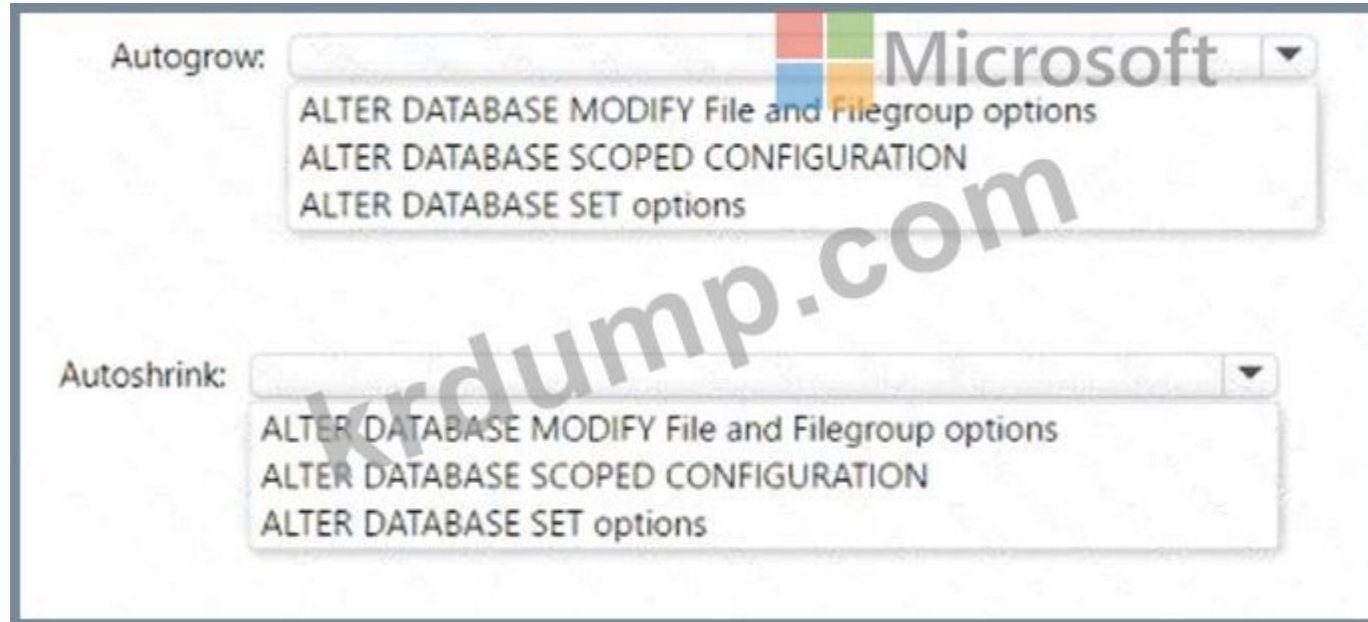
Answer: C (LEAVE A REPLY)

Reference:

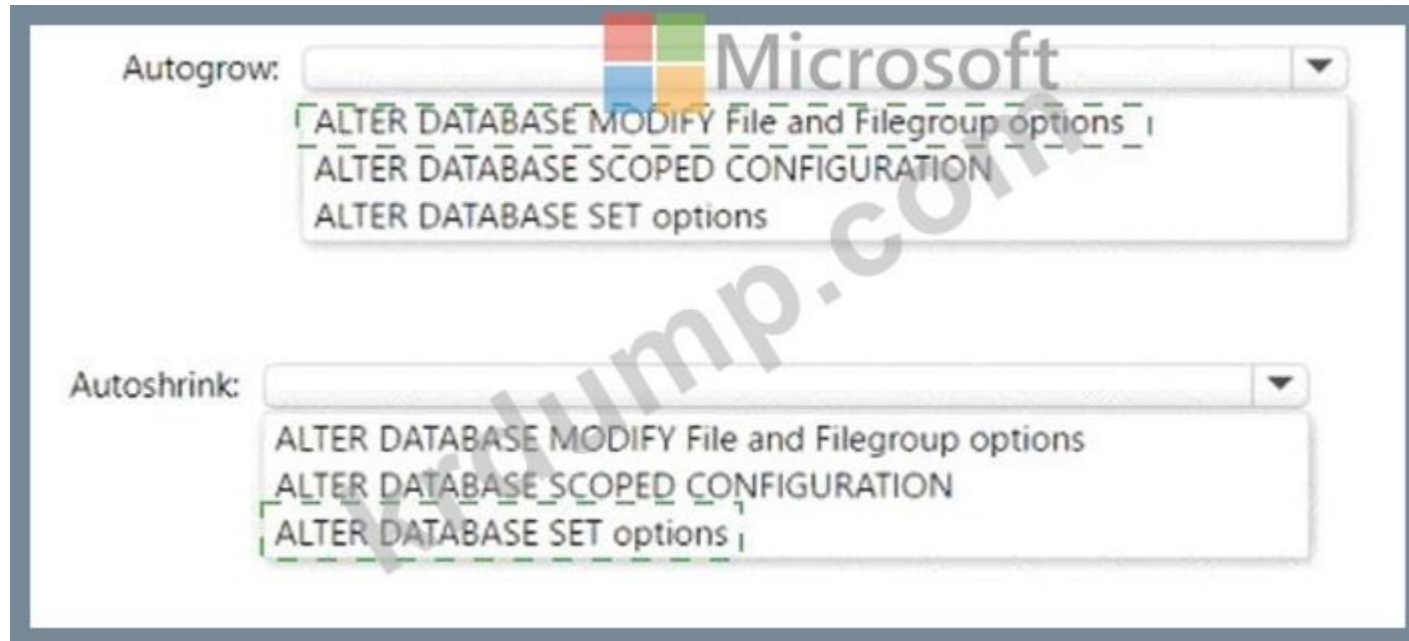
<https://docs.microsoft.com/en-us/azure/dms/dms-overview>

NEW QUESTION: 108

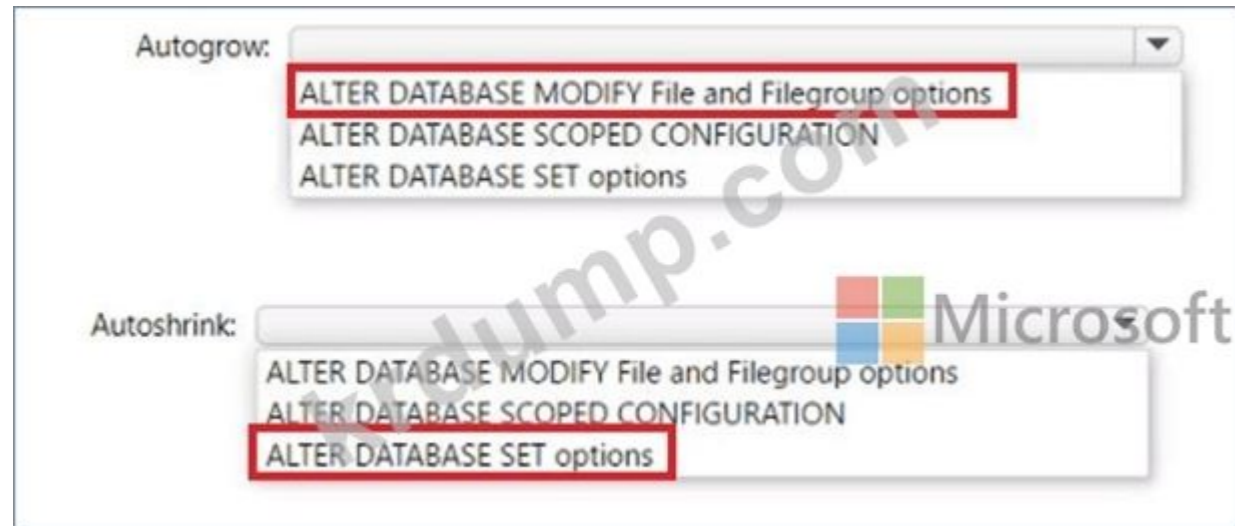
Azure SQL Server Db1 has a database named DB1. The database is currently in the state of `ALTER DATABASE SET options`. You need to ensure that the database is in the state of `ALTER DATABASE MODIFY File and Filegroup options`. What should you do?



Answer:



Explanation:



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

NEW QUESTION: 109

□□ □□□□ □□□□ □□□ □□□□□ □□ □□□ □□□ □□□□ □□□?

- A. □□□□
- B. □□□□ □□
- C. □□
- D. □□□

Answer: (SHOW ANSWER)

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.

Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.

The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/clusters>

NEW QUESTION: 111

10

SQL injection attacks exploit a vulnerability in the application code that interacts with the database. An attacker can inject malicious SQL statements into the user input, such as a form field or a URL parameter, and execute them on the database server, resulting in data theft, corruption, or unauthorized access.

Answer:

See the explanation part for the complete Solution.

Explanation:

SQL injection attacks are a type of cyberattack that exploit a vulnerability in the application code that interacts with the database. An attacker can inject malicious SQL statements into the user input, such as a form field or a URL parameter, and execute them on the database server, resulting in data theft, corruption, or unauthorized access.

To protect all the databases on sql37006S95 from SQL injection attacks, you need to follow some best practices for securing your application and database layers. Here are some of the recommended steps:

- * Use parameterized queries or stored procedures to separate the SQL code from the user input. This will prevent the user input from being interpreted as part of the SQL statement and avoid SQL injection.
- * Validate and sanitize the user input before passing it to the database. This will ensure that the input conforms to the expected format and type, and remove any potentially harmful characters or keywords.
- * Implement least privilege access for the database users and roles. This will limit the permissions and actions that the application can perform on the database, and reduce the impact of a successful SQL injection attack.
- * Enable Advanced Threat Protection for Azure SQL Database. This is a feature that detects and alerts you of anomalous activities and potential threats on your database, such as SQL injection, brute force attacks, or unusual access patterns. You can configure the alert settings and notifications using the Azure portal or PowerShell.

These are some of the steps to protect all the databases on sql37006S95 from SQL injection attacks.

NEW QUESTION: 112

Microsoft SQL Server on Server1 is connected to a database named DB1. User1 is connected to the database. User1 is using the Azure Data Studio to connect to the database. User1 is unable to connect to the database.

* DB1 is an Azure SQL Database.

* Azure Database Migration Service is used to migrate data from Server1 to DB1.

* DB1 is an Azure SQL Database.

* User1 is using the Azure Data Studio to connect to the database.

User1 is unable to connect to the database. User1 is using the Azure Data Studio to connect to the database. User1 is unable to connect to the database.

*: 1000000.

Answer Area



Role:

- Contributor
- SQL DB Contributor
- SQL Server Contributor
- SQL Security Manager

Resource provider:

- Microsoft.Databox
- Microsoft.DataMigration
- Microsoft.Migrate

Answer:

Answer Area



Role:

- Contributor
- SQL DB Contributor
- SQL Server Contributor
- SQL Security Manager

Resource provider:

- Microsoft.Databox
- Microsoft.DataMigration
- Microsoft.Migrate

Explanation:



NEW QUESTION: 113

Q: A company has a data lake in Azure Data Lake Storage. The data lake contains data from various sources, including Azure Synapse Analytics, Azure Databricks, and other data sources. The company wants to ensure that the data in the data lake is secure and compliant with regulatory requirements. Which of the following actions should the company take to ensure data security and compliance?

A.

Actions

- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.
- Create a contained database user.
- Create a resource pool.

Answer Area

Microsoft

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

Answer:

Actions

- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.
- Create a contained database user.
- Create a resource pool.

Answer Area

- Create a resource pool.
- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.

Microsoft

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

Explanation:

- Create a resource pool.
- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15>

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-defined-function?view=sql-server-ver15>

NEW QUESTION: 117

Azure VM SQLVM1 Azure SQL Server User1 SQLVM1 DB1 .
User1 DB1 .
* .
* .
* .
 .
User1 ., .?
: 1.



Answer:



Explanation:



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

NEW QUESTION: 118

Azure Db1 SQL Server .
Db1 .

□□□ □□□ □□□□ □□□? □□□□, □□□□ □□□ □□ □□□□□.

□□: □□ □□□ 1□□□□.

ALTER DATABASE [Db1]

```
Microsoft
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
```

GO

ALTER DATABASE [Db1]

```
Microsoft
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
```

GO

Answer:

ALTER DATABASE [Db1]

```
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
```

GO

ALTER DATABASE [Db1]

```
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
```

GO

Explanation:

The screenshot shows two SQL query windows in Microsoft SQL Server Enterprise Manager. The top window is titled 'ALTER DATABASE [Db1]' and contains the following T-SQL commands:

```
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
```

The bottom window is also titled 'ALTER DATABASE [Db1]' and contains the same list of commands followed by a 'GO' statement:

```
GO
ALTER DATABASE [Db1]
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
GO
```

Box 1: SET AUTOMATIC_TUNING = AUTO

To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:

```
ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO
```

Setting automatic tuning to AUTO will apply Azure Defaults.

Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:

ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON) Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option.

Setting it to OFF will also override any setting that database inherited and disable the tuning option.

Reference:

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

NEW QUESTION: 119

Which SQL statement creates a user in Azure SQL Database that authenticates via an Azure AD application?

A. CREATE USER [app1] FROM EXTERNAL PROVIDER FOR LOGIN app1

B. CREATE USER [app1] FROM EXTERNAL PROVIDER FOR LOGIN app1 WITHOUT LOGIN

C. CREATE USER [app1] FROM EXTERNAL PROVIDER

Answer Area



Answer:

Answer Area



Explanation:

Answer Area



[https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?](https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql)

[view=azuresql](https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql)

NEW QUESTION: 120

Which Azure SQL service supports automatic tuning?

A. Azure SQL Database

B. Azure SQL Data Warehouse

C. Azure SQL Managed Instance

D. Azure SQL Analytics

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

Table1 is an external table in Azure Synapse Analytics. It is created using the following T-SQL script. The script uses the Azure Data Lake Storage Gen2 as the data source. The table has two columns: a primary key column of type datetime and a column of type int. The table is created with the following schema:

A.

B.

Answer: (SHOW ANSWER)

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: 123

Instance1 is a Microsoft SQL Server instance. The instance is configured with the following properties:

```
{
  "type": "Microsoft.Sql/managedInstances",
  "apiVersion": "2019-06-01-preview",
  "name": "[parameters('instanceName')]",
  "location": "[resourceGroup().location]",
  "sku": {
    "name": "[parameters('skuName')]",
    "tier": "[parameters('skuEdition')]"
  },
  "dependsOn": [
    "Microsoft.Resources/deployments/BuildNetworking"
  ],
  "properties": {
    "administratorLogin": "[parameters('adminLogin')]",
    "administratorLoginPassword": "[parameters('adminPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',parameters('netName'),parameters('subnetName'))]"
  }
}
```

licenseType is set to LicenseIncluded and storageRedundancy is set to GeoRedundantStorage. The instance is configured with the following properties:

licenseType is set to LicenseIncluded and storageRedundancy is set to GeoRedundantStorage. The instance is configured with the following properties:



Answer:

Answer Area



Explanation:



NEW QUESTION: 124

□□□ Azure Data Lake Storage Gen2 □□□□□ □□□□ Azure Data Factory □□□□□□ □□□□□ □□□□.

□□ □□□ □□□□ □□□□□ □□□□?

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

Answer: B (LEAVE A REPLY)

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: 125

□□ 1

db1□□□ Azure SQL □□□□□□□□ SalesLT.SalesOrderHeader □□□□ PK_SalesOrderHeader_SalesOrderID □□□□□ □□□□ □□ □□□ □□□ □□□□□ □□□□.

Answer:

See the explanation part for the complete Solution.

Explanation:

To enable page compression on the PK_SalesOrderHeader_SalesOrderID clustered index of the SalesLT.

SalesOrderHeader table in db1, you can use the following Transact-SQL script:

```
-- Connect to the Azure SQL database named db1
```


□□: □□ □□□ 1□□□□.

```
Microsoft
▼ 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://contosol.dfs.core.windows.net/
  applications/year=*//*.parquet',
  FORMAT = 'PARQUET') AS [r]
GO
```

Answer:

```
Microsoft
▼ 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://contosol.dfs.core.windows.net/
  applications/year=*//*.parquet',
  FORMAT = 'PARQUET') AS [r]
GO
```

Explanation:

- A. `ALTER DATABASE DB1 SET`
- B. `ALTER DATABASE DB1 SET`
- C. `ALTER DATABASE AES_256 SET`
- D. `'GeneralPurpose' SET, 'GP_Gen5_2' SET`

Answer: (SHOW ANSWER)

NEW QUESTION: 128

8
 db1 is configured with automatic plan correction. You want to prevent db1 from reverting to the last known good query plan. How do you do this?

Answer:

See the explanation part for the complete Solution.

Explanation:

To prevent db1 from reverting to the last known good query plan, you need to disable the automatic plan correction feature for the database. This feature is enabled by default and allows the Query Store to detect and fix plan performance regressions by forcing the last good plan. However, if you want to test the performance of different plans without interference from the Query Store, you can turn off this feature by using the ALTER DATABASE SCOPED CONFIGURATION statement.

Here are the steps to disable the automatic plan correction feature for db1:

- * Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.
- * Open a new query window and run the following command: ALTER DATABASE SCOPED CONFIGURATION SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = OFF); GO
- * This command will disable the automatic plan correction feature for db1 and allow the Query Optimizer to choose the best plan based on the current statistics and parameters.
- * To verify that the automatic plan correction feature is disabled for db1, you can query the sys.database_scoped_configurations catalog view. The value of the force_last_good_plan column should be 0 for db1.

These are the steps to disable the automatic plan correction feature for db1.

NEW QUESTION: 129

Azure SQL database 25 GB. You want to backup the database using Transact-SQL. How do you do this?

- A. Azure CLI
- B. Azure PowerShell
- C. Azure Portal
- D. Azure Automation

Answer: D (LEAVE A REPLY)

NEW QUESTION: 130

databasebackups folder in the Azure SQL database DB1. You want to backup the database using Transact-SQL. How do you do this? You want to backup the database to the databasebackups folder in the Azure SQL database DB1. How do you do this?

Answer Area



[https://databasebackups.blob.core.windows.net/Backups]

WITH IDENTITY =

'SHARED ACCESS SIGNATURE'	▼
'DatabaseBackups'	
'KeyVault1'	
'SHARED ACCESS SIGNATURE'	

SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-02T19:30:08Z&spr=https&sv=2021-06-08&sr=b&sig=B%2FxEYQi0C%4BqyYCeQwH5z2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'

BACKUP DATABASE DB1

TO URL =

'https://databasebackups.blob.core.windows.net/Backups/db1.bak'

WITH

COPY_ONLY	▼
CHECKSUM	
COMPRESSION	
COPY_ONLY	
DIFFERENTIAL	

Answer:

Answer Area

```
CREATE CREDENTIAL
```

```
[https://databasebackups.blob.core.windows.net/Backups]
```

```
WITH IDENTITY =
```

'SHARED ACCESS SIGNATURE'	▼
'DatabaseBackups'	
'KeyVault1'	
'SHARED ACCESS SIGNATURE'	

```
SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-
```

```
02T19:30:08Z&spr=https&sv=2021-06-
```

```
08&sr=b&sig=B%2FxEYQ%4BqyYCeqlWHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
```

```
BACKUP DATABASE DB1
```

```
TO URL =
```

```
'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
```

```
WITH
```

COPY_ONLY	▼
CHECKSUM	
COMPRESSION	
COPY_ONLY	
DIFFERENTIAL	



krdump.com

Explanation:

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE' ,
SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-
02T19:30:08Z&spr=https&sv=2021-06-
08&sr=b&sig=B%2FxEYQI0C%4BqyYCeQWHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL =
'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY ONLY
```

NEW QUESTION: 131

Two databases, DB1 and DB2, are on a server. The server is configured to use a shared access signature. The server is configured to use a shared access signature. The server is configured to use a shared access signature. The server is configured to use a shared access signature.

- A. The server is configured to use a shared access signature.
- B. The server is configured to use a shared access signature.
- C. The server is configured to use a shared access signature.
- D. The server is configured to use a shared access signature.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 132

VM1 is on an Azure VM. The VM is running Windows Server 2022. The VM is running Windows Server 2022. The VM is running Windows Server 2022. The VM is running Windows Server 2022.

- A. sp_change_users_login
- B. xp_instance_regwrite
- C. xp_grant_login
- D. sp_eddremotelogin

Answer: B (LEAVE A REPLY)

NEW QUESTION: 133

Azure DB1 SQL Server Azure SQL Server Azure

DB1 (HADR)

* (RTO) (RPO)

A. Always On

B.

C. Always On (FCI)

D. Azure

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 134

Azure SQL Server Azure DB1 DB1

* Azure Storage

*

Transact-SQL

: 1

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
MAX_DISPATCH_LATENCY = 3 SEC
```

Answer:

```
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

```
MAX_DISPATCH_LATENCY = 3 SEC)
```

ALLOW_MULTIPLE_EVENT_LOSS

ALLOW_SINGLE_EVENT_LOSS

NO_EVENT_LOSS

Explanation:

Answer Area



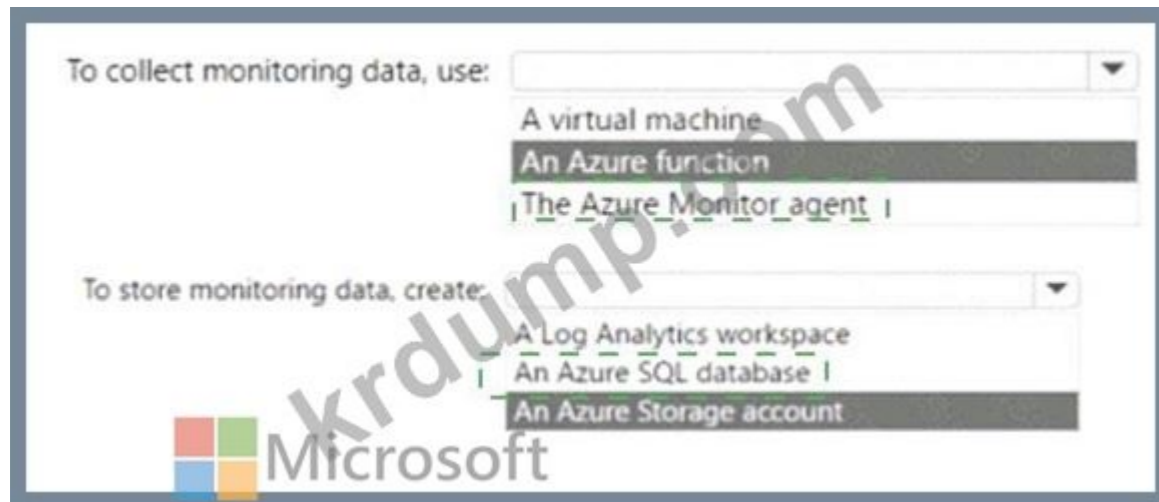
```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
  ACTION (sqlserver.sql_text)
  WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0. event_file
(
  SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'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

NEW QUESTION: 135

□□ □□□ □□□ DB 1□□□ □□□ Azure SQL □□□□□□ □□□□.
SQL Insights□ □□□□ DB 1□ □□□□□□ □□□.
□□□□ □□□ □□□□ □□□? □□□□□ □□□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□□.



Answer:



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: 136

□□ □□□ □□ Azure SQL □□□□□□ □□ □□□ □□□□□.

Point-in-time-restore

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

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

14

Long-term retention

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

Weekly LTR Backups

Keep weekly backups for:

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?

□□ □□ □□ □□ □□□ □ □□ □□□ □□ □□□ □□□ □□□□.
□□: □□ □□ 1□□□□.

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

Answer:

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

Explanation:

Answer Area

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

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

Reference:

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

NEW QUESTION: 140

VM1 is an Azure VM running Windows Server 2022. SQL1 is a Microsoft SQL Server 2019 instance running on VM1. SQL1 is configured to use the Windows Authentication mode. Which of the following stored procedures can be used to connect to SQL1 from a remote machine?

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

Answer: C (LEAVE A REPLY)

NEW QUESTION: 141

DB1 is an Azure SQL Database. Table1 is a table in DB1 with a column of type CHAR(400). The column contains 20 rows of data. The total length of the data in the column is 150,000 characters. Which of the following data types can be used to store the data in the column?

- A. nvarchar(MAX)
- B. nvarchar(MAX)
- C. varchar(MAX)
- D. varchar(200)

Answer: (SHOW ANSWER)

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: 142

Table1 is a table in an Azure Synapse Analytics SQL pool. The table contains 100 rows of data. The total length of the data in the table is 1,000,000 characters. Which of the following data types can be used to store the data in the table?

- A. nvarchar(MAX)
- B. nvarchar(MAX)
- C. varchar(MAX)
- D. varchar(200)

Table1 is a table in an Azure Synapse Analytics SQL pool. The table contains 100 rows of data. The total length of the data in the table is 1,000,000 characters. Which of the following data types can be used to store the data in the table?

- A. nvarchar(MAX)
- B. nvarchar(MAX)

Answer: A (LEAVE A REPLY)

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: 143

DB1 is an Azure SQL Database. You need to create an external table in DB1 that references data in an Azure Storage account. The external table must be created in a schema that is not the default schema. Which command should you use to create the external table?

- A. `CREATE EXTERNAL TABLE`
- B. `CREATE EXTERNAL DATA SOURCE`
- C. `CREATE EXTERNAL SCHEMA`
- D. `CREATE EXTERNAL TABLE WITH SCHEMA`

Answer: (SHOW ANSWER)

NEW QUESTION: 144

RG1 is an Azure Resource Group. You need to create a firewall rule in RG1 that allows traffic to an Azure SQL Server. The rule must be created in a resource group that is not the default resource group. Which command should you use to create the rule?

```
az sql server firewall-rule create \
  --resource-group rg1 \
  --server server1 \
  --name AllowAzureServices \
  --outbound-rule-fqdn allowedFQDN \
  --start-ip-address 0.0.0.0 --end-ip-address 0.0.0.0 \
  --start-ip-address 13.86.217.0 --end-ip-address 13.86.217.255 \
  --start-ipv6-address 9a41:a145:2a80:6c8d:4628:a1b3:5812:3283 --end-ipv6-address 9a41:a145:2a80:6c8d:4628:a1b3:5812:3283
```

Answer:



Explanation:



NEW QUESTION: 145

DB1 is a Microsoft SQL Server 2019 instance on an Azure SQL Managed Instance. The instance is named SQLMI1. The instance is located in the West US region. The instance is currently in a state where it is not accepting connections. You need to ensure that the instance is available and accepting connections. Which of the following actions should you take?

- A. Restart the instance.
- B. Stop and start the instance.
- C. Restart the instance and then start the instance.
- D. Stop the instance.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 146

You are configuring an Azure SQL Managed Instance. You need to ensure that the instance is available and accepting connections. Which of the following actions should you take?

Specify how long you want to keep your point-in-time backups. [Learn more](#)
How many days would you like PITR backups to be kept?

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

Weekly LTR Backups
Keep weekly backups for:

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

Yearly LTR Backups
Keep an annual backup for:

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

□□ □□□□ □□□ □□□□ □□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□□□.

Answer Area

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

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

- 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.

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

Answer:
Answer Area

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

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

- 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.

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

Explanation:

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice] point-in-time restore (PITR) backup.

After the 52nd weekly backup runs, there will be [answer choice] in long term retention. 65 backup copies

NEW QUESTION: 147

Which of the following data types is not supported by the Employee table? (Select all that apply.)

```
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
)
```

Options:

- A. LastHireDate
- B. PhoneNumber
- C. Salary
- D. PhoneNumber
- E. LastHireDate

- A. LastHireDate datetime2(7) NOT NULL
- B. PhoneNumber bigint NOT NULL
- C. Salary Money NOT NULL
- D. PhoneNumber int NOT NULL
- E. LastHireDate datetime2(7) NOT NULL

Answer: C,E (LEAVE A REPLY)

NEW QUESTION: 148

Which of the following methods can be used to migrate data from Microsoft SQL Server (SQL1) to Azure SQL Database (SQLM11)? (Select all that apply.)

- A. Azure Migrate
- B. SQL Server Data Tools (SSDT)
- C. Data Migration Assistant (DMA)

NEW QUESTION: 154

You have a Microsoft SQL Server database on an Azure virtual machine. The database size is 100GB. You need to migrate the database to Azure SQL Database. What should you do?

- A. Use the Azure SQL Database Migration Assistant
- B. Use the Azure SQL Database Migration Assistant
- C. Use the Azure SQL Database Migration Assistant
- D. Use the Azure SQL Database Migration Assistant

Answer: A (LEAVE A REPLY)

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: 155

You have a Microsoft SQL Server database on an Azure virtual machine. The database size is 200GB. You need to migrate the database to Azure SQL Database. What should you do?

* Use the Azure SQL Database Migration Assistant

* Use the Azure SQL Database Migration Assistant, SQL Server Agent, and SQL Server Agent to migrate the database. What should you do?

Answer: Use the Azure SQL Database Migration Assistant.



Answer:

SELECT *

FROM

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

WHERE database_name = 'db1' AND

start_time >

(day, -7, GETDATE())

DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

ORDER BY start_time DESC;

Answer:

SELECT *

FROM

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

WHERE database_name = 'db1' AND

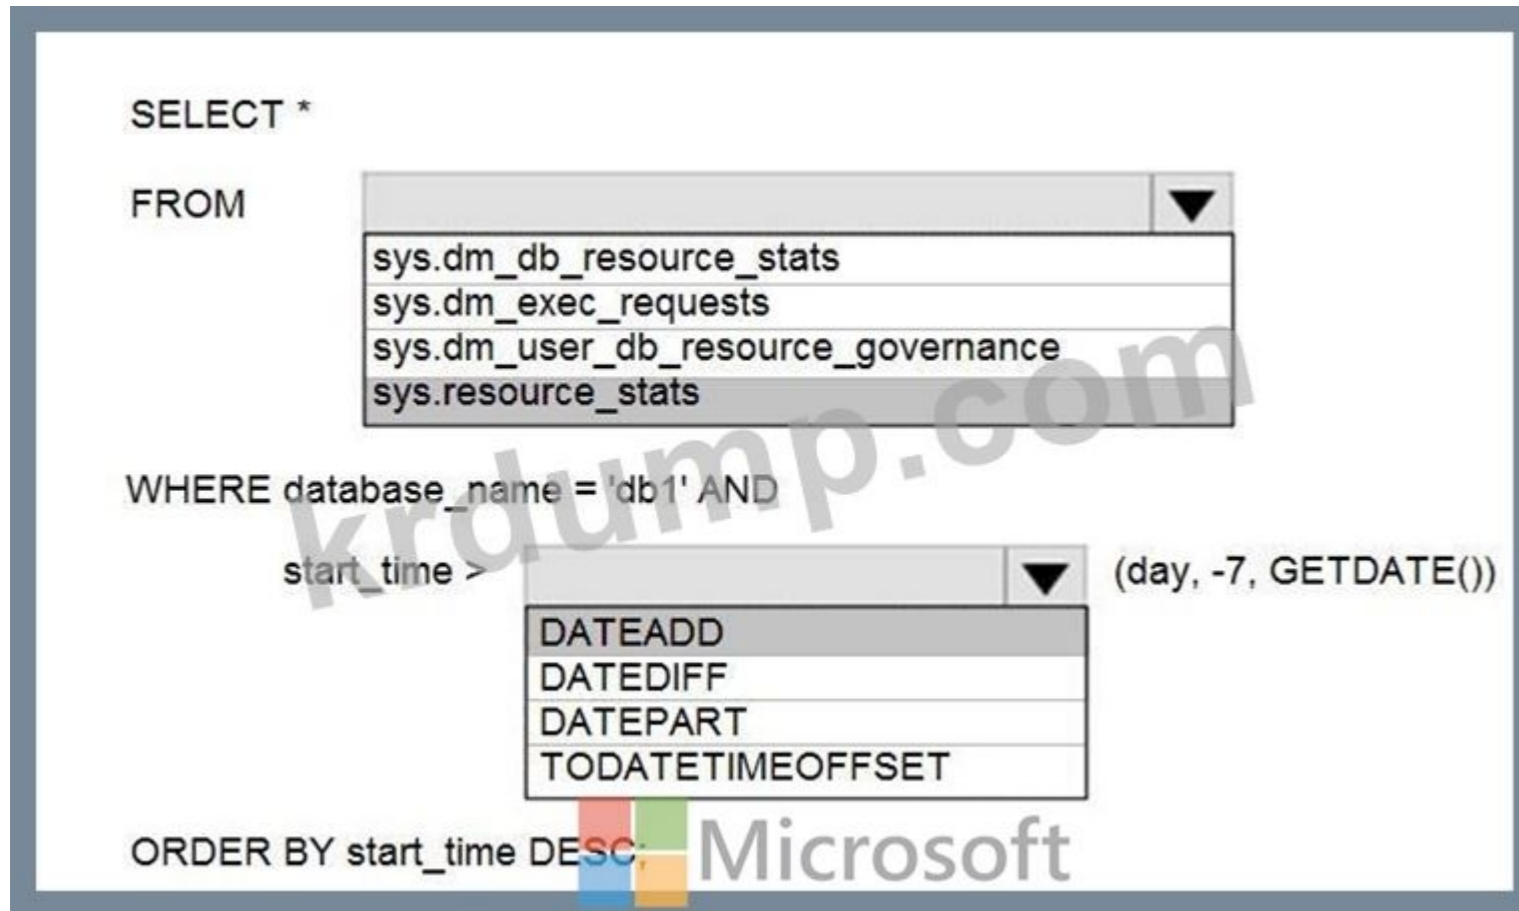
start_time >

(day, -7, GETDATE())

DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

ORDER BY start_time DESC;

Explanation:



Box 1: sys.resource_stats

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

Box 2: DateAdd

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

```
DECLARE @s datetime;
```

```
DECLARE @e datetime;
```

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

```
SET @e= GETUTCDATE();
```

```
SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats WHERE start_time BETWEEN @s AND @e GROUP BY database_name HAVING AVG(avg_cpu_percent) > = 80 Reference:
```

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

DP-300-KR ☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop ☐☐ ☐☐☐☐ ☐☐☐ DP-300-KR ☐☐! DumpTop ☐ ☐☐ **DP-300-KR** ☐☐ ☐☐☐ ☐☐☐☐☐☐☐, DumpTop DP-300-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐☐ ☐☐☐ ☐☐☐☐☐. ☐☐☐☐☐ ☐☐☐☐☐☐☐ ☐☐ DumpTop DP-300-KR ☐☐☐☐☐☐☐. <https://www.dumpst.com/Microsoft/DP-300-KR-dump.html> (410 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)