

# Microsoft.DP-300.v2024-01-23.q244

Exam Code:	DP-300
Exam Name:	Administering Relational Databases on Microsoft Azure
Certification Provider:	Microsoft
Free Question Number:	244
Version:	v2024-01-23
# of views:	858
# of Questions views:	2440
<a href="https://www.freeqas.com/qa/Microsoft/DP-300/Microsoft.DP-300.v2024-01-23.q244.html">https://www.freeqas.com/qa/Microsoft/DP-300/Microsoft.DP-300.v2024-01-23.q244.html</a>	

## NEW QUESTION: 1

You have An Azure SQL managed instance.

You need to configure the SQL Server Agent service to email job notifications.

Which statement should you execute?

- A. 

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

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

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

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

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

## NEW QUESTION: 2

You have an Azure SQL database.

You have a query and the associated execution plan as shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

The performance issue stems from the [answer choice] operator.



	▼
Select	
Index Seek	
Key Lookup	
Nested Loops	

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

	▼
heap	
clustered index	
nonclustered index	

**Answer:**

The performance issue stems from the [answer choice] operator.

	▼
Select	
Index Seek	
Key Lookup	
Nested Loops	

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



	▼
heap	
clustered index	
nonclustered index	

**NEW QUESTION: 3**

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.

You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.

You need to ensure that webapp1 can identify the changes to Table1. The solution must meet the following requirements:

- \* Minimize compute times.
- \* Minimize storage.

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

```

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**

- 1 Connect to SQLDb1 and run the following Transact-SQL statement.  
ALTER DATABASE SQLDb1 SET CHANGE\_TRACKING = ON
- 2 Connect to SQLDb1 and run the following Transact-SQL statement.  
ALTER TABLE dbo.Table1 ENABLE CHANGE\_TRACKING
- 3 From webapp1, connect to SQLDb1, obtain the initial dataset, and run the CHANGETABLE() function.

**Answer:**

**Answer Area**

```

ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON

ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING

From webapp1, connect to SQLDb1....

```

- 1 - ALTER DATABASE SQLDb1 SET CHANGE\_TRACKING = ON
- 2 - ALTER TABLE dbo.Table1 ENABLE CHANGE\_TRACKING
- 3 - From webapp1, connect to SQLDb1....

**NEW QUESTION: 4**

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

You enable columnstore archival compression.

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

NOTE: Each correct selection is worth one point.

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

**Answer: 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.

Plan and Implement a High Availability and Disaster Recovery (HADR) Environment Testlet 1

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office is separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

#### Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

#### Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com.

All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

#### Database Environment

The sales department has the following database workload:

- \* An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

- \* A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

- \* An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1. Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### Current Problems

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

#### Requirements

#### Planned Changes

Litware plans to implement the following changes:

- \* Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

- \* Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

- \* Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

- \* Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

- \* Migrate the SERVER1 databases to the Azure SQL Database platform.

## Technical Requirements

Litware identifies the following technical requirements:

- \* Maintenance tasks must be automated.
- \* The 30 new databases must scale automatically.
- \* The use of an on-premises infrastructure must be minimized.
- \* Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- \* All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

## Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- \* Store encryption keys in Azure Key Vault.
- \* Retain backups of the PII data for two months.
- \* Encrypt the PII data at rest, in transit, and in use.
- \* Use the principle of least privilege whenever possible.
- \* Authenticate database users by using Active Directory credentials.
- \* Protect Azure SQL Database instances by using database-level firewall rules.
- \* Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

## Business Requirements

Litware identifies the following business requirements:

- \* Meet an SLA of 99.99% availability for all Azure deployments.
- \* Minimize downtime during the migration of the SERVER1 databases.
- \* Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- \* Once all requirements are met, minimize costs whenever possible.

### NEW QUESTION: 5

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.

Which switch should you use to switch between languages?

- A. Error! Hyperlink reference not valid.]
- B. %<language>
- C. Error! Hyperlink reference not valid.]
- D. @<language>

**Answer: B (LEAVE A REPLY)**

You can override the default language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

Reference:

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

### NEW QUESTION: 6

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named App1. App1 uses db1.

You need to enable Resource Governor for a App1. The solution must meet the following requirements:

App1 must be able to consume all available CPU resources.

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

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

**Actions**

- Create a plan.
- Create a classifier function in db1.
- Create a workload group.
- Create a classifier function in the master database.
- Create a resource pool that has the following configurations.  
MAX\_CPU\_PERCENT = 100  
MIN\_CPU\_PERCENT = 50

**Answer Area**

**Answer:**

### Answer Area

- 1 - Create a resource pool that has the following configurations...
- 2 - Create a workload group.
- 3 - Create a classifier function in the master database.

### NEW QUESTION: 7

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

You plan to deploy an Always On availability group that will have the following configurations: VM1 will host the primary replica of DB1.

VM2 will host a secondary replica of DB1.

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

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

DATABASE MyDB1

FROM DISK = 'D:\DB1.bak'

WITH

GO

Microsoft

**Answer:**

DATABASE MyDB1

FROM DISK = 'D:\DB1.bak'

WITH

GO

Microsoft

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server?view=sql-server-ver15>

### NEW QUESTION: 8

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

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

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

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

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

Answer:

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

Explanation

Enable Database Mail.
Create an operator.
Add a failure notification to the job.

Text Description automatically generated

Reference:

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

### NEW QUESTION: 9

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



You need to reduce the Log 10 percentage. The solution must minimize costs.

What should you do?

- A. B. Change RecoverymodeltoSimple.
- B. Increase the number of vCores.
- C. Performcheckpoint operation.
- D. ChangeService tier toBusiness Critical.


**Answer:** ([SHOW ANSWER](#))


### NEW QUESTION: 10

You create all of the tables and views for ResearchDB1.

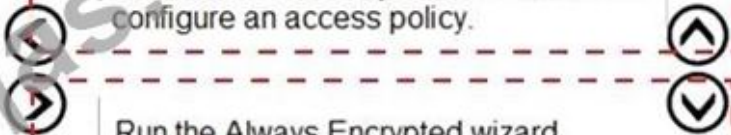
You need to implement security for ResearchDB1. The solution must meet the security and compliance requirements.


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

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



**Answer:**

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



Explanation

Graphical user interface, text, application Description automatically generated

Register ResearchApp1 to Azure AD.

Create an Azure Key Vault instance and configure an access policy.

Run the Always Encrypted wizard.



Reference:

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

**NEW QUESTION: 11**

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

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

NOTE: Each correct selection is worth one point.

Azure Database Migration Service pricing tier: Standard 2-vCore, Standard 4-vCore, Premium 4-vCore

Required Azure resource: A virtual network that has service endpoints, A VPN gateway, An Azure Logic app

**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

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: 12**

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

What should you create?

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

**Answer: B (LEAVE A REPLY)**

Scenario: Contoso requirements for the sales transaction dataset include:

Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the

table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

Reference:

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

**NEW QUESTION: 13**

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

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

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

NOTE: Each correct selection is worth one point.



**Answer:**



Explanation

Answer Area



**NEW QUESTION: 14**

You have an Azure SQL database named DB 1 in the General Purpose service tier.

You need to monitor DB 1 by using SQL Insights.

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

NOTE: Each correct selection is worth one point.

To collect monitoring data, use:

- A virtual machine
- An Azure function**
- The Azure Monitor agent

To store monitoring data, create:

- A Log Analytics workspace
- An Azure SQL database
- An Azure Storage account**

Answer:

To collect monitoring data, use:

- A virtual machine
- An Azure function**
- The Azure Monitor agent**

To store monitoring data, create:

- A Log Analytics workspace
- An Azure SQL database**
- An Azure Storage account**

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: 15**

You have an Azure data factory that has two pipelines named PipelineA and PipelineB.

PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



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

Operator: Greater than

Aggregation type: Total

Threshold value: 2

Aggregation granularity (Period): 5 minutes

Frequency of evaluation: Every 5 minutes

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

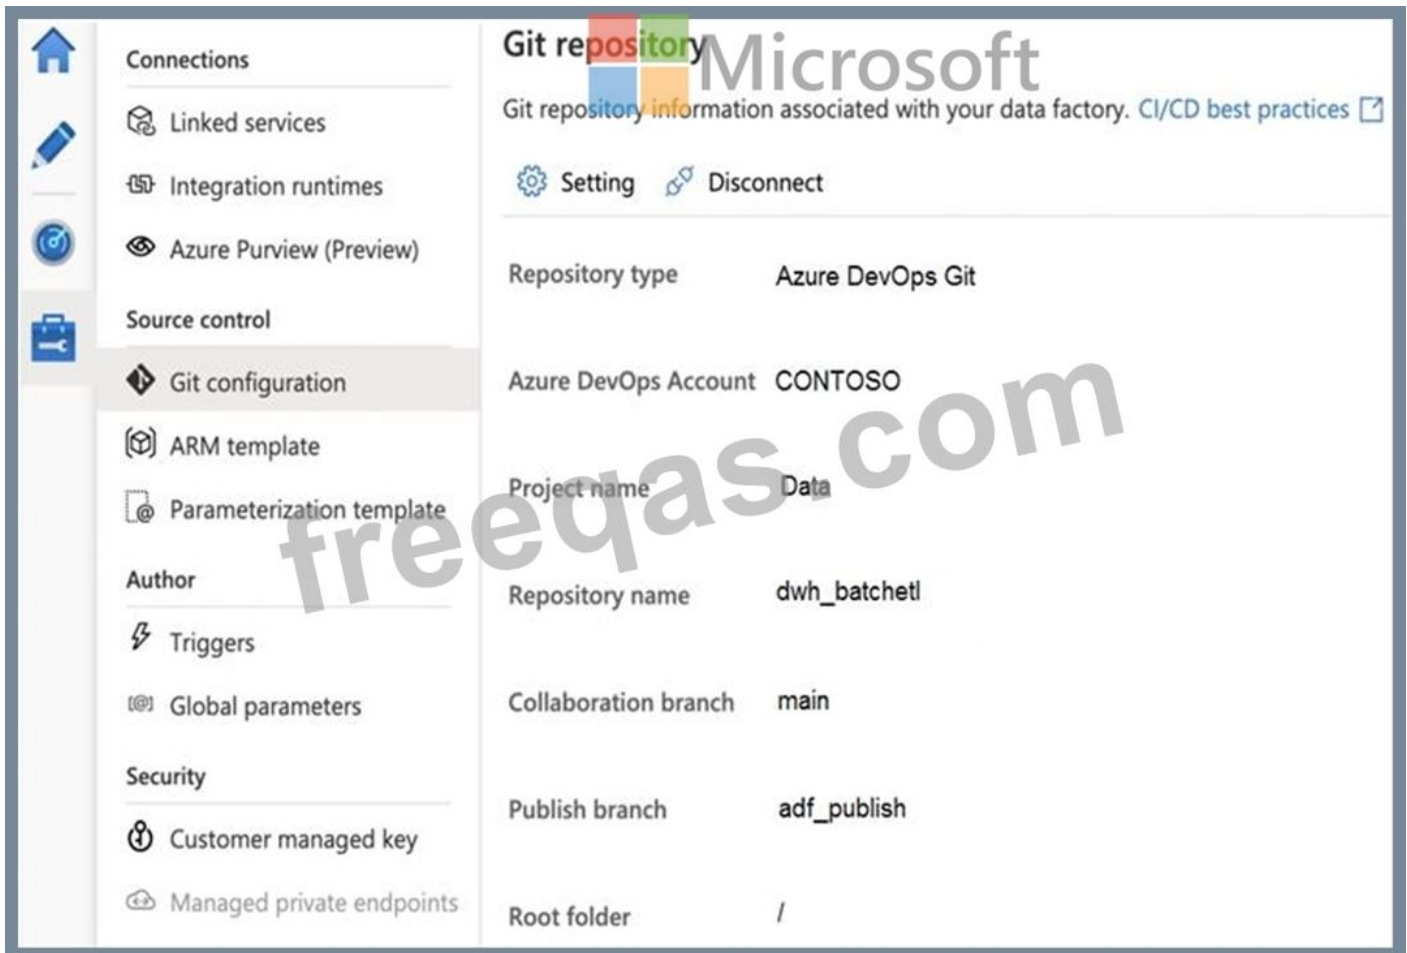
Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input checked="" type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input checked="" type="radio"/>
An alert notification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

Reference:

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

### NEW QUESTION: 16

You configure version control for an Azure Data Factory instance as shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

- /
- adf\_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh\_batchetl/adf\_publish/contososales
- /main

**Answer:**

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

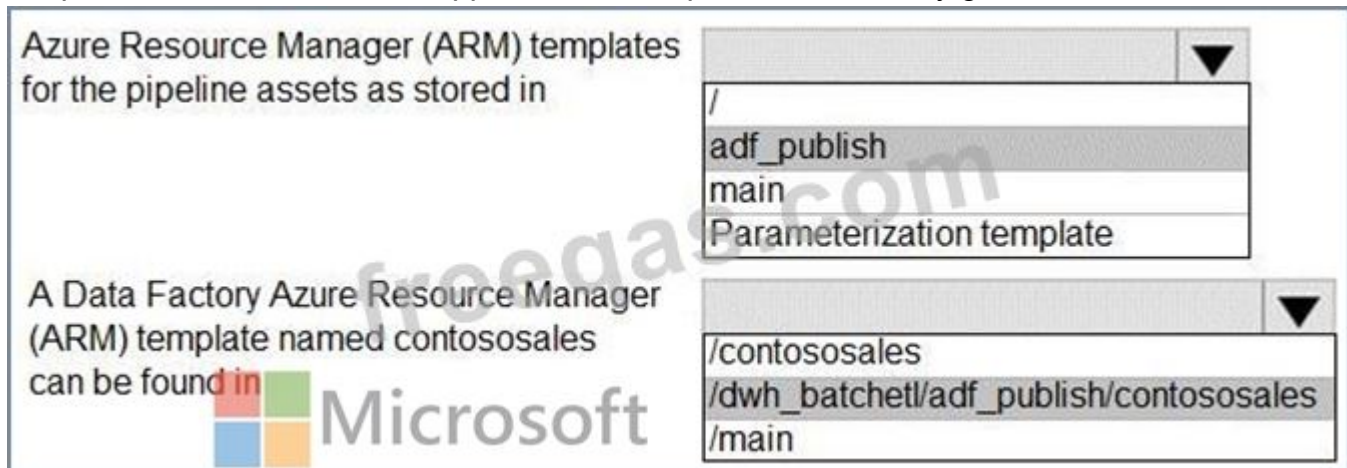
- /
- adf\_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh\_batchetl/adf\_publish/contososales
- /main

## Explanation

Graphical user interface, text, application Description automatically generated



Box 1: adf\_publish

By default, data factory generates the Resource Manager templates of the published factory and saves them into a branch called adf\_publish. To configure a custom publish branch, add a publish\_config.json file to the root folder in the collaboration branch. When publishing, ADF reads this file, looks for the field publishBranch, and saves all Resource Manager templates to the specified location. If the branch doesn't exist, data factory will automatically create it. An example of what this file looks like is below:

```
{  
  "publishBranch": "factory/adf_publish"  
}
```

Box 2: /dwh\_barchlet/ adf\_publish/contososales

RepositoryName: Your Azure Repos code repository name. Azure Repos projects contain Git repositories to manage your source code as your project grows. You can create a new repository or use an existing repository that's already in your project.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/source-control>

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

## NEW QUESTION: 17

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

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

NOTE: Each correct selection is worth one point.

Table type to store retail store data: [Dropdown with options: Hash, Replicated, Round-robin]

Table type to store promotional data: [Dropdown with options: Hash, Replicated, Round-robin]

**Answer:**

Table type to store retail store data: [Dropdown with options: Hash, Replicated, Round-robin]

Table type to store promotional data: [Dropdown with options: Hash, Replicated, Round-robin]

Reference:

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

### NEW QUESTION: 18

You have 20 Azure SQL databases provisioned by using the vCore purchasing model. You plan to create an Azure SQL Database elastic pool and add the 20 databases. Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. total size of all the databases
- B. geo-replication support
- C. number of concurrently peaking databases \* peak CPU utilization per database
- D. maximum number of concurrent sessions for all the databases
- E. total number of databases \* average CPU utilization per database

**Answer: A,C,E (LEAVE A REPLY)**

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model:  $\text{MAX}(\text{Total number of DBs} \times \text{average vCore utilization per DB}, \text{Number of concurrently peaking DBs} \times \text{Peak vCore utilization per DB})$

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

Reference:

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

### NEW QUESTION: 19

You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance requirements.

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

**Actions**

- Create a user in the master database
- Modify the Azure SQL server administrator account
- Create contained database users
- Create an Azure AD administrator for the logical server
- Connect to the databases by using an Azure AD account
- Enable the contained database authentication option

**Answer Area**

Navigation arrows: > (right), < (left), ^ (up), v (down)

**Answer:**

**Actions**


- Create a user in the master database
- Modify the Azure SQL server administrator account
- Create contained database users
- Create an Azure AD administrator for the logical server
- Connect to the databases by using an Azure AD account
- Enable the contained database authentication option

**Answer Area**

- Create an Azure AD administrator for the logical server
- Create contained database users
- Connect to the databases by using an Azure AD account

Navigation arrows: > (right), < (left), ^ (up), v (down)

Explanation



Scenario: Authenticate database users by using Active Directory credentials.

The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

- \* Create and populate Azure AD.
- \* Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
- \* Create an Azure Active Directory administrator. (Step 1)
- \* Configure your client computers.
- \* Create contained database users in your database mapped to Azure AD identities. (Step 2)
- \* Connect to your database by using Azure AD identities. (Step 3)

Reference:

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

Topic 3, ADatum Corporation

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the button to return to the question.

Overview

ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.

#### Existing Environment

ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.

SALESDB collects data from the stores and the website.

DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.

REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.

#### Requirements

##### Planned Changes

ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:

- \* Migrate SALESDB and REPORTINGDB to an Azure SQL database.
- \* Migrate DOCDB to Azure Cosmos DB.
- \* 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.
- \* As they arrive, all the sales documents in JSON format must be transformed into one consistent format.
- \* Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.

##### Technical Requirements

The new Azure data infrastructure must meet the following technical requirements:

- \* Data in SALESDB must be encrypted by using Transparent Data Encryption (TDE). The encryption must use your own key.
- \* SALESDB must be restorable to any given minute within the past three weeks.
- \* Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.
- \* Missing indexes must be created automatically for REPORTINGDB.
- \* Disk IO, CPU, and memory usage must be monitored for SALESDB.

#### **NEW QUESTION: 20**

You have an Azure SQL managed instance.

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

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

NOTE: Each correct selection is worth one point.

RESTORE  DB1 FROM

DATABASE	DISK = N'\\NAS01\SQLBackups\DB1.bak';
FILE	TAPE = N'\\Tape0'
LOG	URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

**Answer:**

RESTORE  DB1 FROM

DATABASE	DISK = N'\\NAS01\SQLBackups\DB1.bak';
FILE	TAPE = N'\\Tape0'
LOG	URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

**NEW QUESTION: 21**

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

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

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

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

**Answer:**

Actions	Answer Area
Encrypt the Salary column by using the randomized encryption type.	Create a column master key.
Create a column encryption key.	Create a column encryption key.
Enable Transparent Data Encryption (TDE).	Encrypt the Salary column by using the randomized encryption type.
Encrypt the Salary column by using the deterministic encryption type.	
Apply a dynamic data mask to the Salary column.	
Create a column master key.	

### Explanation

Step 1: Create a column master key

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

Step 2: Create a column encryption key.

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

Randomized encryption uses a method that encrypts data in a less predictable manner.

Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

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

Reference:

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

### NEW QUESTION: 22

You have the following Azure Resource Manager template.

```

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

```



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

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

**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

A screenshot of a computer Description automatically generated with low confidence

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: 23

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

What should you include in the recommendation?

- A. For each customer, grant the customer db\_writerto the existing schema.
- B. For each customer, create an additional schema and grant the customer ddl\_admin to the new schema.
- C. For each customer, grant the customer ddl\_admin to the existing schema.
- D. For each customer, create an additional schema and grant the customer db\_writerto the new schema.

Answer: ([SHOW ANSWER](#))

### NEW QUESTION: 24

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


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

The database fails on Wednesday at 10:00.

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

Actions	Answer Area
Monday, Tuesday, and then Wednesday differential backups	
Wednesday, Tuesday, and then Monday log backups	
full backup	
Monday, Tuesday, and then Wednesday log backups	
Wednesday, Tuesday, and then Monday differential backups	
Wednesday log backups	
Wednesday differential backup	

Navigation icons: > < < >




Answer:

Actions	Answer Area
Monday, Tuesday, and then Wednesday differential backups	
Wednesday, Tuesday, and then Monday log backups	
full backup	
Monday, Tuesday, and then Wednesday log backups	
Wednesday, Tuesday, and then Monday differential backups	
Wednesday log backups	
Wednesday differential backup	

Navigation icons: > < < >

full backup
Wednesday differential backup
Wednesday log backups



**NEW QUESTION: 25**

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

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

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 26**

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

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

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

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

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

Solution: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Does this meet the goal?

- A. Yes
- B. No

**Answer: B (LEAVE A REPLY)**

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: 27**

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

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

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

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

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

Does this meet the goal?

A. Yes

B. No

**Answer: B (LEAVE A REPLY)**

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


Reference:

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

### NEW QUESTION: 28

You have an Azure data factory that has two pipelines named PipelineA and PipelineB.

PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



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

Operator: Greater than

Aggregation type: Total

Threshold value: 2

Aggregation granularity (Period): 5 minutes

Frequency of evaluation: Every 5 minutes

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Reference:

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

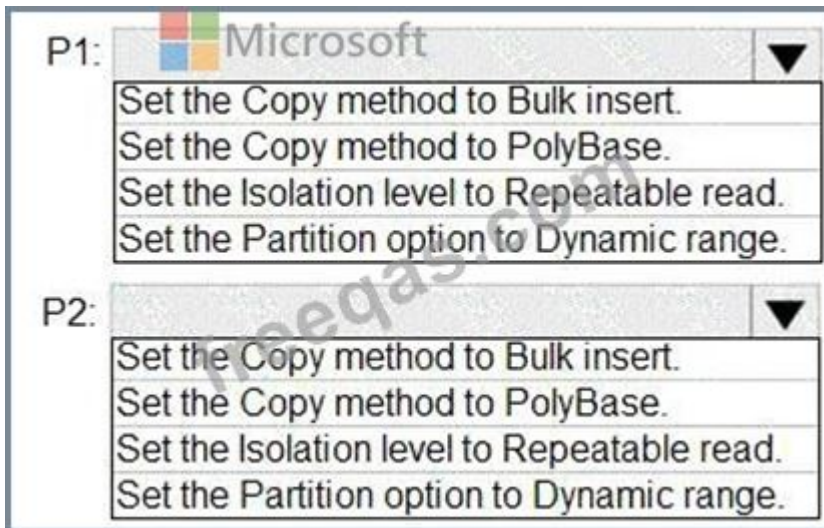
### NEW QUESTION: 29

You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.

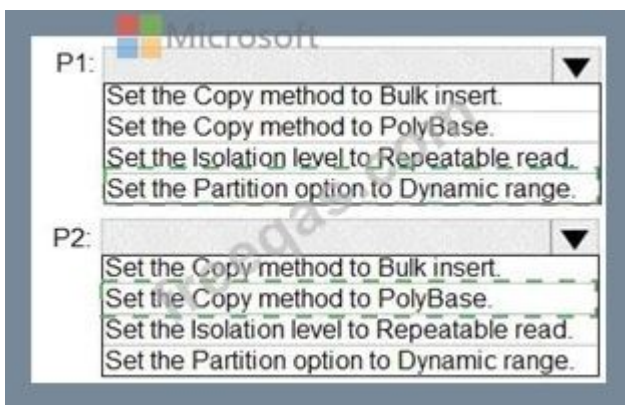
ADF1 contains the following pipelines:

P1: Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account  
 P2: Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2  
 You need to configure P1 and P2 to maximize parallelism and performance.

Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

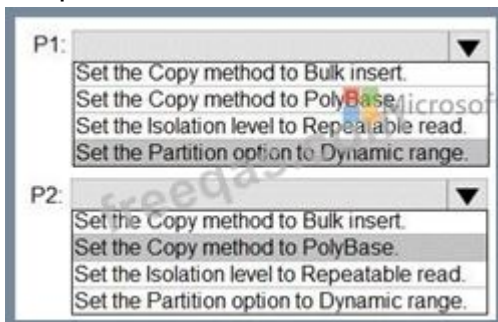


**Answer:**



**Explanation**

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



P1: Set the Partition option to Dynamic Range.

The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel.

P2: Set the Copy method to PolyBase

Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store.

(Polybase supports Azure Blob storage and Azure Data Lake Store by default.) Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse>

<https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse>

**NEW QUESTION: 30**

You have an Azure SQL Database instance named DatabaseA on a server named Server1. You plan to add a new user named App1 to DatabaseA and grant App1 db\_datacenter permissions. App1 will use SQL Server Authentication.

You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.

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

**Actions**

**Answer Area**

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

On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssw0rd!';

On DatabaseA, run ALTER ROLE db\_datareader ADD MEMBER [App1];

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

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



**Answer:**

**Answer Area**

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

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

On DatabaseA run ALTER ROLE db\_datareader ADD Member [App1]

- 1 - On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'
- 2 - On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]
- 3 - On DatabaseA run ALTER ROLE db\_datareader ADD Member [App1]

Reference:

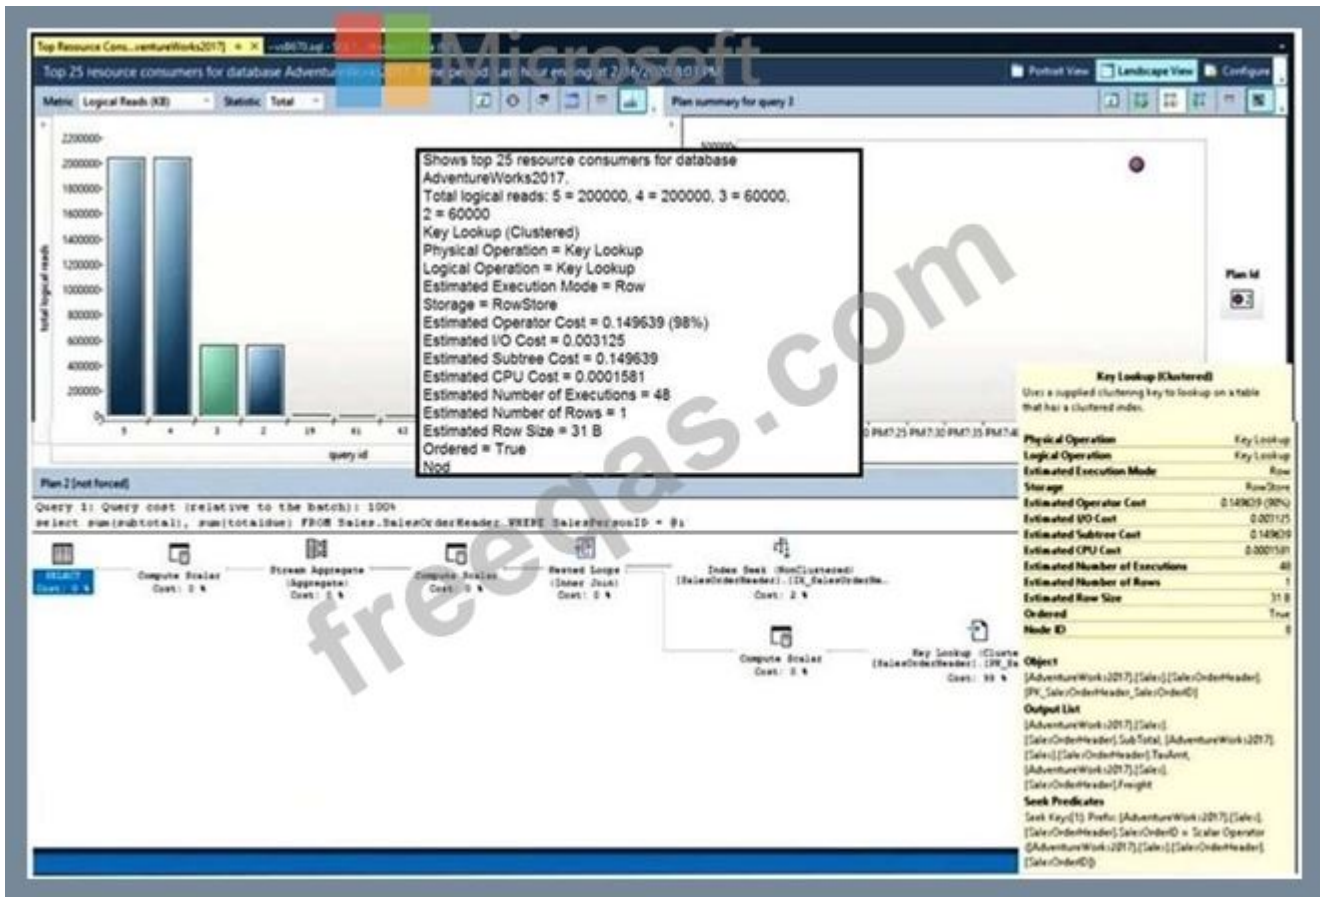
<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

**NEW QUESTION: 31**

**HOTSPOT**

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



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

Hot Area:

Answer Area



Statements

Yes

No

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



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



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



Answer:

## Answer Area

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<input checked="" 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/Reference:

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

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!

PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 32

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements. What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Purchasing model:   
DTU  
vCore

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

**Answer:**

Purchasing model:

	▼
Azure virtual machine reserved instances	
DTU	
vCore	

Deployment option:

	▼
An Azure SQL Database elastic pool	
An Azure SQL Database managed instance	
A SQL Server Always On availability group	



Microsoft

Reference:

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

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

**NEW QUESTION: 33**

HOTSPOT

You have an Azure SQL database named db1.

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area  
 Microsoft

SELECT \*

FROM

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

WHERE database\_name = 'db1' AND

start\_time >

▼
DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

(day, -7, GETDATE())

ORDER BY start\_time DESC;

Answer:

Answer Area

SELECT \*

FROM

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

WHERE database\_name = 'db1' AND

start\_time >

▼
DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

(day, -7, GETDATE())

ORDER BY start\_time DESC;

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
```

Incorrect Answers:

sys.dm\_exec\_requests:  
sys.dm\_exec\_requests returns information about each request that is executing in SQL Server. It does not have a column named database\_name.

sys.dm\_db\_resource\_stats:

sys.dm\_db\_resource\_stats does not have any start\_time column.

Note: sys.dm\_db\_resource\_stats returns CPU, I/O, and memory consumption for an Azure SQL Database database. One row exists for every 15 seconds, even if there is no activity in the database. Historical data is maintained for approximately one hour.

sys.dm\_user\_db\_resource\_governance returns actual configuration and capacity settings used by resource governance mechanisms in the current database or elastic pool. It does not have any start\_time column.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-database>  
Optimize Query Performance Testlet 1 This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem

statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

### Existing Environment

#### Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

#### Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com.

All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

#### Database Environment

The sales department has the following database workload:

- \* An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

- \* A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

- \* An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1 Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### Current Problems

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

#### Requirements

#### Planned Changes

Litware plans to implement the following changes:

- \* Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

- \* Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

- \* Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- \* Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- \* Migrate the SERVER1 databases to the Azure SQL Database platform.

#### Technical Requirements

Litware identifies the following technical requirements:

- \* Maintenance tasks must be automated.
- \* The 30 new databases must scale automatically.
- \* The use of an on-premises infrastructure must be minimized.
- \* Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- \* All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

#### Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- \* Store encryption keys in Azure Key Vault.
- \* Retain backups of the PII data for two months.
- \* Encrypt the PII data at rest, in transit, and in use.
- \* Use the principle of least privilege whenever possible.
- \* Authenticate database users by using Active Directory credentials.
- \* Protect Azure SQL Database instances by using database-level firewall rules.
- \* Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

#### Business Requirements

Litware identifies the following business requirements:

- \* Meet an SLA of 99.99% availability for all Azure deployments.
- \* Minimize downtime during the migration of the SERVER1 databases.
- \* Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- \* Once all requirements are met, minimize costs whenever possible.

### NEW QUESTION: 34

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Configuration
DB1	Azure SQL Database	Hyperscale service tier No secondary replicas
App1	Azure Web Apps	App1 has read-only access to DB1. There are multiple instances of App1.

You need to create a read-only replica of DB1 and configure the App1 instances to use the replica.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

To add read-only replicas of DB1:

- Create a replica on the same logical server.
- Create a new logical server and configure geo-replication.
- Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

- Add an ApplicationIntent entry to the connection string.
- Add a MultiSubnetFailover entry to the App1 connection string.
- Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

**Answer:**

**Answer Area**

To add read-only replicas of DB1:

- Create a replica on the same logical server.
- Create a new logical server and configure geo-replication.
- Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

- Add an ApplicationIntent entry to the connection string.
- Add a MultiSubnetFailover entry to the App1 connection string.
- Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

Reference:

<https://sqlserverguides.com/read-only-replica-azure-sql/>

**NEW QUESTION: 35**

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

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

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

## Actions

## Answer Area

Encrypt the Salary column by using the randomized encryption type.


Create a column encryption key.

Enable Transparent Data Encryption (TDE).

Encrypt the Salary column by using the deterministic encryption type.

Apply a dynamic data mask to the Salary column.

Create a column master key.



Microsoft

## Answer:

Answer Area
Create a column master key.
Create a column encryption key.
Encrypt the Salary column by using the randomized encryption type.

- 1 - Create a column master key
- 2 - Create a column encryption key.
- 3 - Encrypt the Salary column by using the randomized encryption type.

Reference:

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

## NEW QUESTION: 36

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

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

## Actions

## Answer Area

Create and configure a schedule.

Create a SQL Server Agent job.

Publish the runbook.

Create an Azure Automation account.

Import the SqlServer module.

Create a runbook that runs a PowerShell script.

Run `sp_add_jobserver`.



Answer:

**Answer Area**



---

Create an Azure Automation account.

---

Import the SqlServer module.

---

Create a runbook that runs a PowerShell script.

---

Create and configure a schedule.

---

- 1 - Create an Azure Automation account.
- 2 - Import the SqlServer module.
- 3 - Create a runbook that runs a PowerShell script.
- 4 - Create and configure a schedule.

Reference:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statistics-maintenance-using/ba-p/368974>

## Topic 2, Contoso Ltd

### Existing Environment

Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.

### Active Directory

Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active Directory.

### Database Environment

Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

Name	Role
SQL1	Primary data warehouse
SQL2	Secondary data warehouse
SQL3	Extract, transform, and load (ETL) server

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.

The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.

### Current Business Model

Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers. The customers receive the files by using FTP.

### Requirements

#### Planned Changes

Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource group that contains an Azure SQL database will be provisioned for the customer. The database will have a complete copy of the financial data. The data to which each customer will have access will depend on the service agreement tier. The customers can change tiers by changing their service agreement.

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

- ⋮

#### Business Goals

Contoso identifies the following business requirements:

- ⋮

#### Technical Requirements

Contoso identifies the following technical requirements:

- ⋮

#### Monitoring Requirements

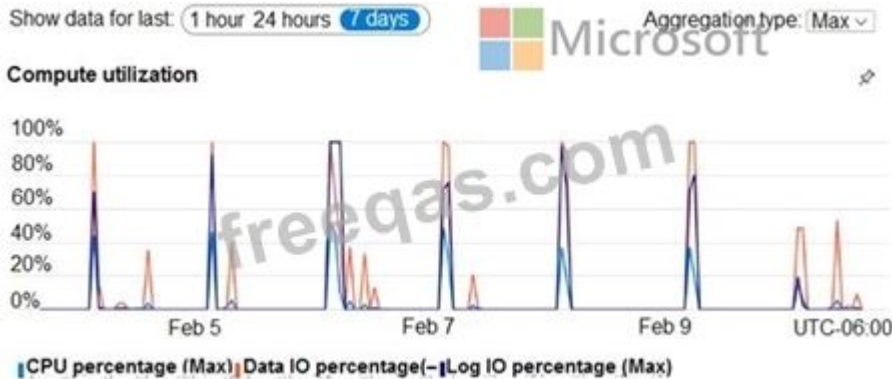
Contoso identifies the following monitoring requirements:

- ⋮

:

## PaaS Prototype

During prototyping of the PaaS solution in Azure, you record the compute utilization of a customer's Azure SQL database as shown in the following exhibit.



## Role Assignments

For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

+ Add Edit columns Refresh Remove Got feedback?

Check access Role assignments Deny assignments Classic administrators Roles

Manage access to Azure resources for users, groups, service principals and managed identities at this scope by creating role assignments. [Learn more](#)

Number of role assignments for this subscription 15 2000

Search by name or email Groups 2 selected All scopes

Group by Role

Showing a filtered set of results. Total number of role assignments: 15

2 items (2 Groups)

Name	Type	Role	Scope
DBA Group 1	Group	Contributor	This resource
DBA Group 2	Group	SQL DB Contributor	This resource

## NEW QUESTION: 37

You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named 081. You need to ensure that Group 1 has read access to new tables created in 061. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.



**Answer:**



**Explanation**



**NEW QUESTION: 38**

You have the following Azure Data Factory pipelines:

- \* Ingest Data from System1
- \* Ingest Data from System2
- \* Populate Dimensions
- \* Populate Facts

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

What should you do to schedule the pipelines for execution?

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

**Answer: D (LEAVE A REPLY)**

Section: [none]

Explanation/Reference:

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

**NEW QUESTION: 39**

You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group.

You need to recommend a deployment solution that meets the following requirements:

- \* Provides a Service Level Agreement (SLA) of at least 99.95%
- \* Replicates databases in the same group synchronously
- \* Minimizes the latency of database writes

What should you recommend?

- A.** Create a proximity group and an availability set. Deploy each virtual machine to the availability set. Add both virtual machines to the proximity group.
- B.** Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group.
- C.** Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group.
- D.** Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group.

**Answer: A (LEAVE A REPLY)**

### NEW QUESTION: 40

You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1. You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.

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

Actions

- Run CREATE LOGIN user1@contoso.com FROM EXTERNAL PROVIDER on the master database.
- Run ALTER SERVER ROLE securityadmin ADD MEMBER user1@contoso.com.
- Create a managed identity for SQLMI1.
- Grant SQLMI1 read access to Azure AD.
- Run CREATE USER user1@contoso.com FROM LOGIN user1@contoso.com.

Answer Area

freedocs.com

**Answer:**

## Answer Area

Grant SQLMI1 read access to Azure AD.

Run CREATE LOGIN,,,,,

Run AKTER SERVER,,,,,

1 - Grant SQLMI1 read access to Azure AD.

2 - Run CREATE LOGIN,,,,,

3 - Run AKTER SERVER,,,,,

Reference:

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

### NEW QUESTION: 41

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.

You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.

You need to ensure that webapp1 can identify the changes to Table1. The solution must meet the following requirements:

- \* Minimize compute times.
- \* Minimize storage.

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

#### Actions

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:**

**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**

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

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

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

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

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the change function.

Explanation

Graphical user interface, text, application Description automatically generated

**Answer Area**


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

```
ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON
```
- 2 Connect to SQLDb1 and run the following Transact-SQL statement.  

```
ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING
```
- 3 From webapp1, connect to SQLDb1, obtain the initial dataset, and run the CHANGETABLE() function.

**NEW QUESTION: 42**

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



PipelineB has two activities as shown in the following exhibit.



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

Aggregation type: Total

Threshold value: 2

Aggregation granularity (Period): 5 minutes

Frequency of evaluation: Every 5 minutes

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Reference:

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

### NEW QUESTION: 43

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.

You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

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

**Actions**
**Answer Area**


Deploy Grafana to an Azure virtual machine.

Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.

Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.



**Answer:**

**Answer Area**  


---

Configure your Azure Databricks cluster to use the Databricksmonitoring library.

---



---

Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file

---



---

Create Dropwizard counters in your application code

---

- 1 - Configure your Azure Databricks cluster to use the Databricksmonitoring library.
- 2 - Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file
- 3 - Create Dropwizard counters in your application code

**NEW QUESTION: 44**

What should you use to migrate the PostgreSQL database?

- A. Azure Data Box
- B. AzCopy
- C. Azure Database Migration Service
- D. Azure Site Recovery

**Answer: (SHOW ANSWER)**

Reference:

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

Topic 2, Litware

Existing Environment

Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

#### Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

#### Database Environment

The sales department has the following database workload:

An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1. Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### Current Problems

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

#### Requirements

#### Planned Changes

Litware plans to implement the following changes:

Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.

Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Migrate the SERVER1 databases to the Azure SQL Database platform.

#### Technical Requirements

Litware identifies the following technical requirements:

Maintenance tasks must be automated.

The 30 new databases must scale automatically.

The use of an on-premises infrastructure must be minimized.

Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

#### Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

Store encryption keys in Azure Key Vault.

Retain backups of the PII data for two months.

Encrypt the PII data at rest, in transit, and in use.

Use the principle of least privilege whenever possible.

Authenticate database users by using Active Directory credentials.

Protect Azure SQL Database instances by using database-level firewall rules.

Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

#### Business Requirements

Litware identifies the following business requirements:

Meet an SLA of 99.99% availability for all Azure deployments.

Minimize downtime during the migration of the SERVER1 databases.

Use the Azure Hybrid Use Benefits when migrating workloads to Azure.

Once all requirements are met, minimize costs whenever possible.

### **NEW QUESTION: 45**

#### HOTSPOT

You have an Azure Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

You need to recommend a solution for account1 that meets the following requirements:

- \* Automatically deletes the logs at the end of each retention period
- \* Minimizes storage costs

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area** Microsoft

To minimize storage costs:

- Store the infrastructure logs and the application logs in the Archive access tier.
- Store the infrastructure logs and the application logs in the Cool access tier.
- Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

- Azure Data Factory pipelines
- Azure Blob storage lifecycle management rules
- Immutable Azure Blob storage time-based retention policies

**Answer:**

**Answer Area**



To minimize storage costs:

- Store the infrastructure logs and the application logs in the Archive access tier.
- Store the infrastructure logs and the application logs in the Cool access tier.
- Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

- Azure Data Factory pipelines
- Azure Blob storage lifecycle management rules
- Immutable Azure Blob storage time-based retention policies

Section: [none]

Explanation:

Box 1: Store the infrastructure logs in the Cool access tier the application logs in the Archive access tier Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Azure Blob storage lifecycle management rules

Blob storage lifecycle management offers a rich, rule-based policy that you can use to transition your data to the best access tier and to expire data at the end of its lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**NEW QUESTION: 46**

You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.

What should you include in the plan?

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

**Answer: D (LEAVE A REPLY)**

Reference:

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

### Topic 3, ADatum Corporation

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the button to return to the question.

#### Overview

ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.

#### Existing Environment

ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.

SALESDB collects data from the stores and the website.

DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.

REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.

#### Requirements

#### Planned Changes

ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:

Migrate SALESDB and REPORTINGDB to an Azure SQL database.

Migrate DOCDB to Azure Cosmos DB.

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. As they arrive, all the sales documents in JSON format must be transformed into one consistent format.

Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.

#### Technical Requirements

The new Azure data infrastructure must meet the following technical requirements:

Data in SALESDB must be encrypted by using Transparent Data Encryption (TDE). The encryption must use your own key.

SALESDB must be restorable to any given minute within the past three weeks.

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

Missing indexes must be created automatically for REPORTINGDB.

Disk IO, CPU, and memory usage must be monitored for SALESDB.

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!

PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps,

**40%OFF Special Discount: Exam-Tests**)

#### **NEW QUESTION: 47**

You have an Azure subscription that contains an Azure SQL database.

The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource\_semaphore waits.

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

NOTE: Each correct selection is worth one point.

SELECT

Microsoft  
is\_user\_process  
wait\_time  
wait\_type

SUM(wait\_time) AS total\_wait\_time\_ms

FROM sys.

dm\_exec\_query\_stats  
dm\_exec\_requests  
query\_store\_query

JOIN sys.dm\_exec\_sessions AS dmv2

ON dmv1.session\_id = dmv2.session\_id

WHERE is\_user\_process = 1

GROUP BY wait\_type

ORDER BY SUM(wait\_time) DESC;

Answer:

```

SELECT
    is_user_process
    wait_time
    wait_type
    SUM(wait_time) AS total_wait_time_ms
FROM sys.
    dm_exec_query_stats
    dm_exec_requests
    query_store_query
JOIN sys.dm_exec_sessions AS dmvs2
    ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;

```

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

**NEW QUESTION: 48**

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

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

NOTE: Each correct selection is worth one point.

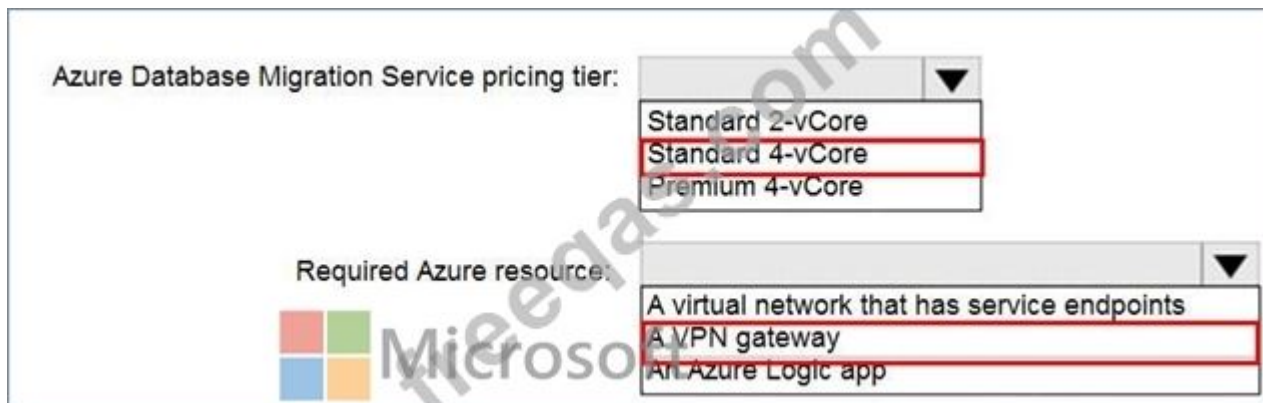
Azure Database Migration Service pricing tier:

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

Required Azure resource:

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

**Answer:**



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: 49

You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database.

Which database option should set?

- A. PARAMETERIZATIONto FORCED
- B. PARAMETERIZATIONto SIMPLE
- C. Delayed Durability to Forced
- D. READ\_COMMITTED\_SNAPSHOTto ON

**Answer: (SHOW ANSWER)**

In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:

\* The READ COMMITTED isolation level with the READ\_COMMITTED\_SNAPSHOT database option set to ON.

\* The SNAPSHOT isolation level.

If READ\_COMMITTED\_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Incorrect Answers:

A: When the PARAMETERIZATION database option is set to SIMPLE, the SQL Server query optimizer may choose to parameterize the queries. This means that any literal values that are contained in a query are substituted with parameters. This process is referred to as simple parameterization. When SIMPLE parameterization is in effect, you cannot control which queries are parameterized and which queries are not.

B: You can specify that all queries in a database be parameterized by setting the PARAMETERIZATION database option to FORCED. This process is referred to as forced parameterization.

C: Delayed transaction durability is accomplished using asynchronous log writes to disk. Transaction log records are kept in a buffer and written to disk when the buffer fills or a buffer flushing event takes place.

Delayed transaction durability reduces both latency and contention within the system.

Some of the cases in which you could benefit from using delayed transaction durability are:

- \* You can tolerate some data loss.
- \* You are experiencing a bottleneck on transaction log writes.
- \* Your workloads have a high contention rate.

Reference:

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

### NEW QUESTION: 50

You configure version control for an Azure Data Factory instance as shown in the following exhibit.

Property	Value
Repository type	Azure DevOps Git
Azure DevOps Account	CONTOSO
Project name	Data
Repository name	dwh_batchetl
Collaboration branch	main
Publish branch	adf_publish
Root folder	/

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

NOTE: Each correct selection is worth one point.

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

- /
- adf\_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh\_batchetl/adf\_publish/contososales
- /main

**Answer:**

Azure Resource Manager (ARM) templates for the pipeline assets as stored in

- adf\_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh\_batchetl/adf\_publish/contososales
- /main

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/source-control>

**NEW QUESTION: 51**

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

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

You need to create a backup of DB1 that is accessible to the Azure Database Migration Service. What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Run:

- A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH FILE\_SNAPSHOT option

Store the backup in:

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

**Answer:**

Run:

- A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH FILE\_SNAPSHOT option

Store the backup in:

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

Reference:

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

### NEW QUESTION: 52

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A. distributed availability groups
- B. database mirroring
- C. log shipping
- D. Database Migration Assistant

**Answer: D (LEAVE A REPLY)**

Explanation

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

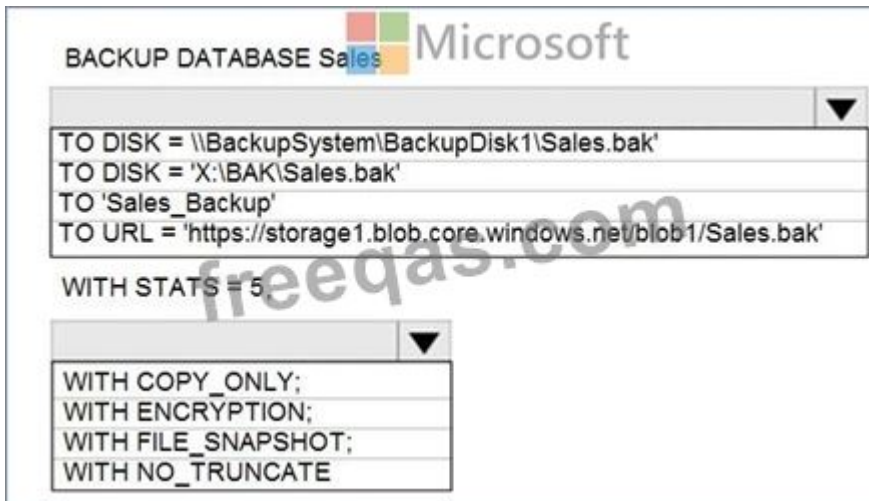
### NEW QUESTION: 53

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales.

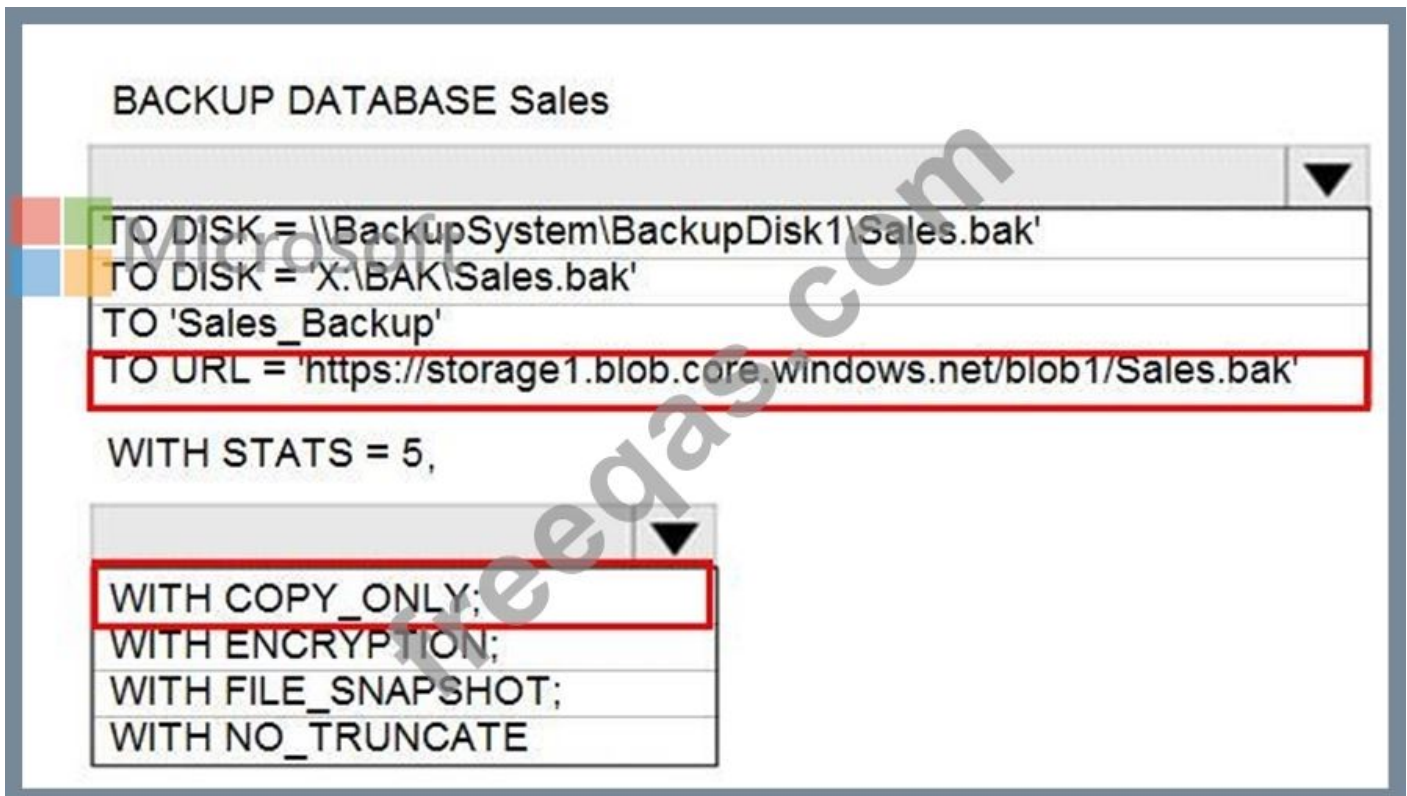
You need to initiate a backup of Sales.

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

NOTE: Each correct selection is worth one point.



**Answer:**



Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-instance/ba-p/386154>

**NEW QUESTION: 54**

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

You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers.

The solution must prevent all the salespeople from viewing or inferring the credit card information.

What should you include in the recommendation?

- A. row-level security
- B. data masking

- C. Always Encrypted
- D. column-level security

**Answer: (SHOW ANSWER)**

Explanation

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

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

Example:

XXXX-XXXX-XXXX-1234

**NEW QUESTION: 55**

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

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

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

NOTE: Each correct selection is worth one point.

```
EXECUTE AS 

|        |   |
|--------|---|
|        | ▼ |
| CALLER |   |
| LOGIN  |   |
| OWNER  |   |
| USER   |   |

 = 'user1@contoso.com
```

GO



```
SELECT SUSER_SNAME ( )
```

	▼
REVERT	
REVOKE	
ROLLBACK	

GO  
**Answer:**

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

CALLER
LOGIN
OWNER
USER

```
GO
```

```
SELECT SUSER_SNAME()
```

REVERT
REVOKE
ROLLBACK



```
GO
```

Reference:

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

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

### NEW QUESTION: 56

You have SQL Server on an Azure virtual machine that contains a database named Db1.

You need to enable automatic tuning for Db1.

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

NOTE: Each correct selection is worth one point.

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

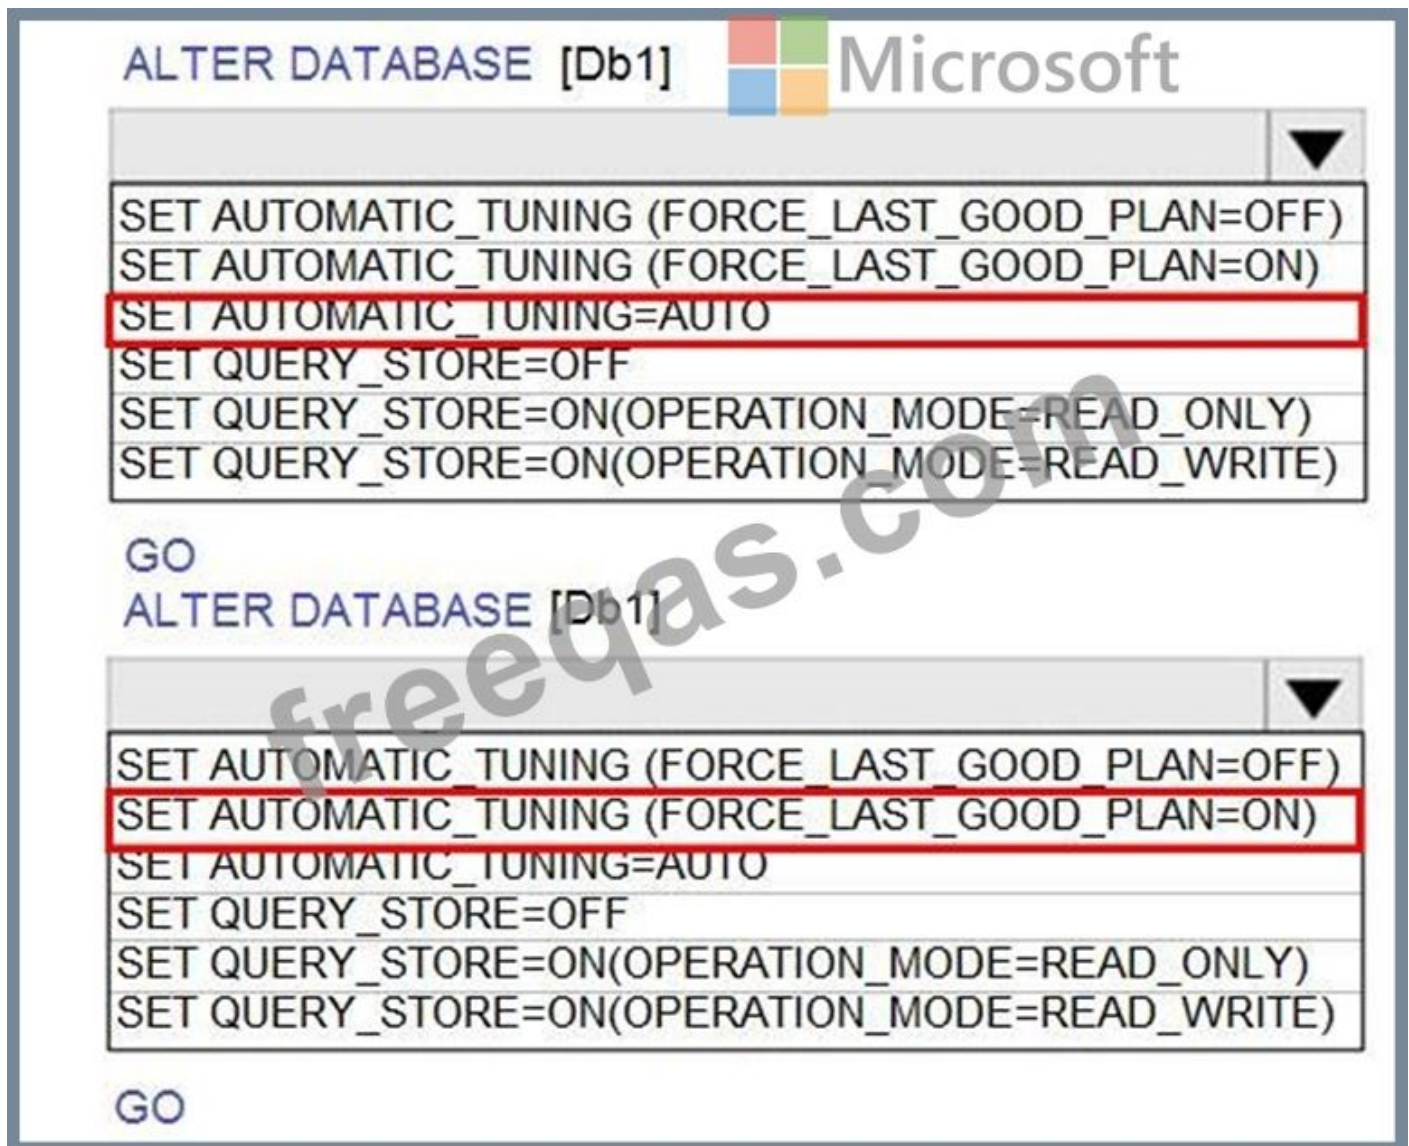
ALTER DATABASE [Db1]

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



Answer:



Reference:

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

### NEW QUESTION: 57

You have an Azure SQL database named sqldb1.

You need to minimize the amount of space by the data and log files of sqldb1.

What should you run?

- A. DBCC SHRINKDATABASE
- B. sp\_clean\_db\_free\_space
- C. sp\_clean\_db\_file\_free\_space
- D. DBCC SHRINKFILE

**Answer: A (LEAVE A REPLY)**

Explanation

DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql>

**NEW QUESTION: 58**

You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1.

Which isolation level should you configure?

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

**Answer:** ([SHOW ANSWER](#))

Reference:

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

**NEW QUESTION: 59**

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state.

What should you do?

- A. Double the value of Data Flush interval
- B. Decrease by half the value of Data Flush Interval
- C. Double the value of Statistics Collection Interval
- D. Decrease by half the value of Statistics Collection interval

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

Explanation

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the maximum size limit between storage size checks, it transitions to read-only mode.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

**NEW QUESTION: 60**

You have a SQL Server on Azure Virtual Machines instance named VM1 that hosts a database named DB1. You run the following query.

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

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

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

**Answer:**

Answer Area

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

**NEW QUESTION: 61**

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

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

**Answer: D (LEAVE A REPLY)**


Explanation

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

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

Timeline Description automatically generated

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



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

Reference:

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

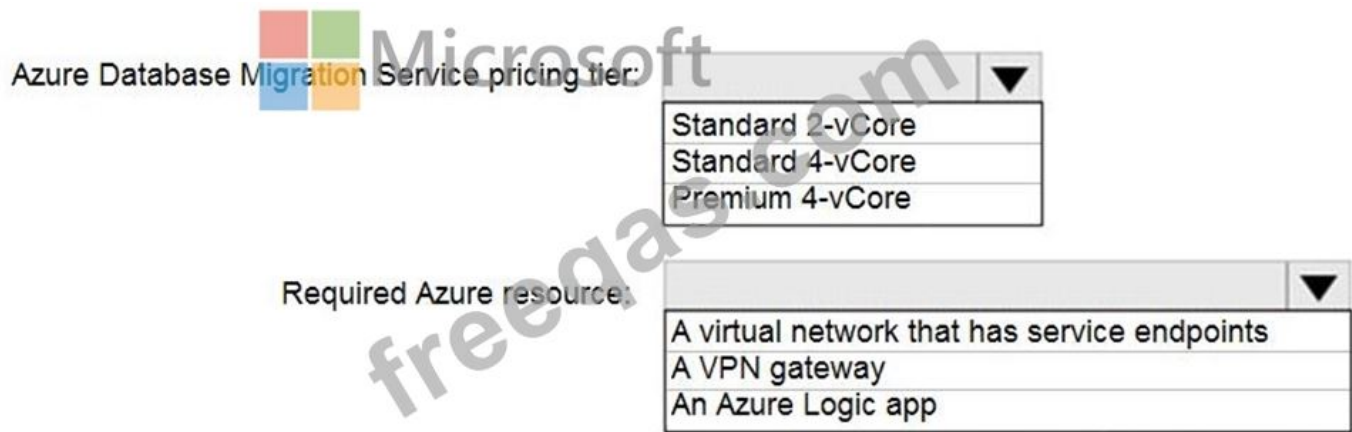
**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 62

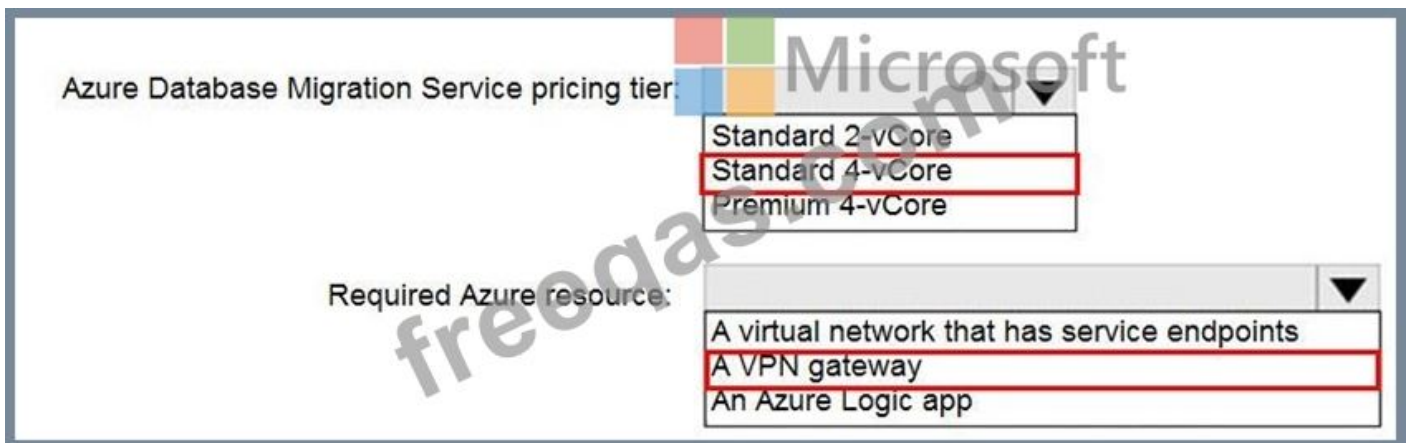
You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.

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

NOTE: Each correct selection is worth one point.



Answer:



Explanation:

Azure Database Migration service

Box 1: Premium 4-VCORE

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

Minimize downtime during the migration of the SERVER1 databases.

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

Incorrect Answers:

The Standard pricing tier suits most small- to medium- business workloads, but it supports offline migration only.

Box 2: A VPN gateway

You need to create a Microsoft Azure Virtual Network for the Azure Database Migration Service by using the Azure Resource Manager deployment model, which provides site-to-site connectivity to your on-premises source servers by using either ExpressRoute or VPN.

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: 63

You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2.

You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable query optimizer fixes.
- B. Enable Resource Governor.
- C. Enable parameter sniffing.
- D. Create a workload group.
- E. Configure In-memory OLTP.
- F. Run the Database Engine Tuning Advisor.
- G. Reduce the Max Degree of Parallelism value.

**Answer: B,C (LEAVE A REPLY)**

Reference:

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

### NEW QUESTION: 64

You have a SQL Server on Azure Virtual Machines instance named VM1 that hosts a database named DB1. You run the following query.

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

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

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

**Answer:**

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

### NEW QUESTION: 65

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

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

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

**Answer Area**

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

**Answer:**

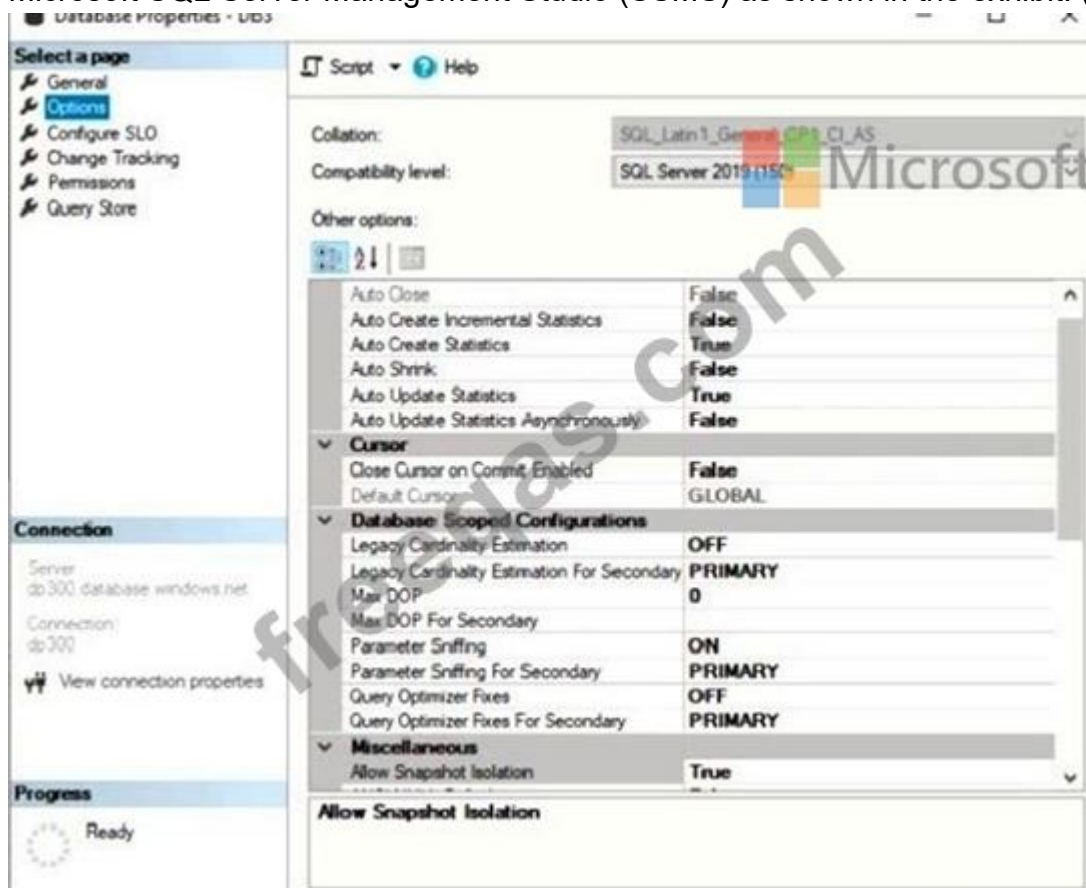
**Answer Area**

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

### NEW QUESTION: 66

You have an Azure SQL database named DB3.

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



Which Transact-SQL command should you run?

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

**Answer: C (LEAVE A REPLY)**

The exhibits displays Database [State] properties.

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

Reference:

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

### NEW QUESTION: 67

You have an Azure SQL database named db1.

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

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

NOTE: Each correct selection is worth one point.

SELECT \*

FROM

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

WHERE database\_name = 'db1' AND

start\_time >

	▼
DATEADD	
DATEDIFF	
DATEPART	
TODATETIMEOFFSET	

(day, -7, GETDATE())

ORDER BY start\_time DESC;

**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())
ORDER BY start_time DESC;

```

Reference:

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

**NEW QUESTION: 68**

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

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

NOTE: Each correct selection is worth one point.

Table type to store retail store data:

- Hash
- Replicated
- Round-robin

Table type to store promotional data:

- Hash
- Replicated
- Round-robin

**Answer:**

Table type to store retail store data:

- Hash
- Replicated
- Round-robin

Table type to store promotional data:

- Hash
- Replicated
- Round-robin

Reference:

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

**NEW QUESTION: 69**

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

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

What should you use?

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

**Answer: A (LEAVE A REPLY)**

Explanation

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

Reference:

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

**NEW QUESTION: 70**

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.

You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

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

Deploy Grafana to an Azure virtual machine.

Build a `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file.

Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.

**Answer:**

## ACTIONS

Deploy Grafana to an Azure virtual machine.

Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.

Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.

Explanation

Graphical user interface, text, application Description automatically generated with medium confidence

Configure the Databricks cluster to use the Databricks monitoring library.

Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.

Create Dropwizard counters in the application code.

Send application metrics using Dropwizard.

Spark uses a configurable metrics system based on the Dropwizard Metrics Library.

To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps:

Step 1: Configure your Azure Databricks cluster to use the Databricks monitoring library.

Prerequisite: Configure your Azure Databricks cluster to use the monitoring library.

Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file Step 3: Create

Dropwizard counters in your application code Create Dropwizard gauges or counters in your application code

## NEW QUESTION: 71

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

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

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

**ACTIONS**

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

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

Create a managed identity for SQLMI1.

Grant SQLMI1 read access to Azure AD.

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

**Answer Area****Answer:****Answer Area**

Grant SQLMI1 read access to Azure AD.

Run CREATE LOGIN user1@contoso.com....

Run ALTER SERVER ROLE securityadmin...

- 1 - Grant SQLMI1 read access to Azure AD.
- 2 - Run CREATE LOGIN user1@contoso.com....
- 3 - Run ALTER SERVER ROLE securityadmin...

Reference:

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

**NEW QUESTION: 72**

What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?

- A. Configure SalesSQLDb1 as writable.
- B. Update the connection strings of SalesSQLDb1App1.
- C. Update the firewall rules of SalesSQLDb1.
- D. Update the users in SalesSQLDb1.

**Answer: C (LEAVE A REPLY)**

Explanation

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

**NEW QUESTION: 73**

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.

You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

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

**Actions****Answer Area**

Deploy Grafana to an Azure virtual machine.

Build a `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file.


Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.

⊖⊕

⬆⬇



**Answer:**

**Answer Area**

Configure your Azure Databricks cluster to use the Databricks monitoring library.

Build the `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file

1 - Configure your Azure Databricks cluster to use the Databricks monitoring library.

2 - Build the `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file

### NEW QUESTION: 74

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Azure region
VM1	Azure virtual machine	West US 2
MI1	Azure SQL Managed Instance	East US

You need to configure a connection between VM1 and MI1. The solution must meet the following requirements:

- \* The connection must be encrypted.
- \* Network latency must be minimized.

What should you implement?

- A. private endpoints
- B. virtual network peering
- C. service endpoints
- D. a site-to-site VPN

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

### NEW QUESTION: 75

You have an Azure subscription.

You plan to deploy an Azure SQL database by using an Azure Resource Manager template.

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

NOTE: Each correct selection is worth one point.

#### Answer Area

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



**Answer:**

**Answer Area**

```

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

```

Reference:

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

**NEW QUESTION: 76**

You are evaluating the role assignments.

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

NOTE: Each correct selection is worth one point.

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

**Answer:**

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

Reference:

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

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
 PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

#### NEW QUESTION: 77

You have an Azure SQL database named DB1. You run a query while connected to DB1. You review the actual execution plan for the query, and you add an index to a table referenced by the query.

You need to compare the previous actual execution plan for the query to the Live Query Statistics.

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

- A. For DB1, set QUERY\_CAPTURE\_MODE of Query Store to All.
- B. Run the SET SHOWPLAN\_ALL Transact-SQL statement.
- C. Save the actual execution plan.
- D. Enable Query Store for DB1.

**Answer: C (LEAVE A REPLY)**

Section: [none]

Explanation:

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

This option can compare between:

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

One active execution plan and one previously saved query execution plan.

Two selected query plans in Query Store.

### NEW QUESTION: 78

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

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

**Answer:** ([SHOW ANSWER](#))

Explanation

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

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

Timeline Description automatically generated

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



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

Reference:

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

### NEW QUESTION: 79

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.

Which resource provider should you enable?

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

**Answer: B (LEAVE A REPLY)**

Explanation

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger>

### NEW QUESTION: 80

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

You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1.

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



The screenshot shows a drag-and-drop interface for configuring an Azure SQL managed instance. On the left, under the heading "Actions", there is a list of five actions: "Create a workload group.", "Create a user-defined classifier function.", "Modify Resource Governor.", "Create a contained database user.", and "Create a resource pool." On the right, there is an "Answer Area" which is currently empty. Between the two areas are four circular arrows: two pointing left and two pointing right, indicating the drag-and-drop functionality. A large watermark "freeqs.com" is visible across the interface.

**Answer:**

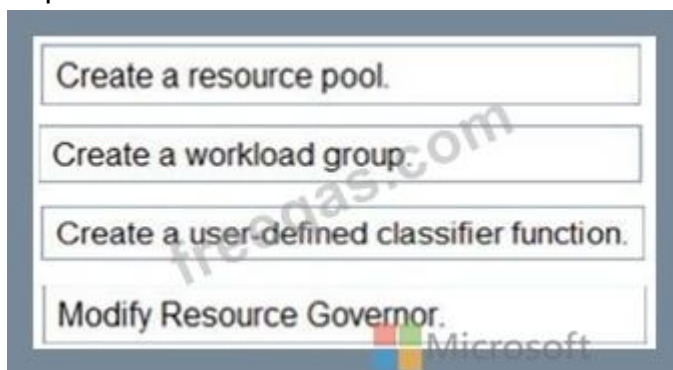
## Actions

- 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.

## Explanation



Text, table Description automatically generated

Reference:

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

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-defi>

## NEW QUESTION: 81

You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database.

Which database option should set?

- A. PARAMETERIZATIONto FORCED
- B. PARAMETERIZATIONto SIMPLE
- C. Delayed Durability to Forced
- D. READ\_COMMITTED\_SNAPSHOTto ON

**Answer: D (LEAVE A REPLY)**

Section: [none]

Explanation:

In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of

uncommitted data modifications using either:

\* The READ COMMITTED isolation level with the READ\_COMMITTED\_SNAPSHOT database option set to ON.

\* The SNAPSHOT isolation level.

If READ\_COMMITTED\_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine

uses row versioning to present each statement with a transactionally consistent snapshot of the data as it

existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Incorrect Answers:

A: When the PARAMETERIZATION database option is set to SIMPLE, the SQL Server query optimizer may

choose to parameterize the queries. This means that any literal values that are contained in a query are

substituted with parameters. This process is referred to as simple parameterization. When SIMPLE

parameterization is in effect, you cannot control which queries are parameterized and which queries are not.

B: You can specify that all queries in a database be parameterized by setting the PARAMETERIZATION

database option to FORCED. This process is referred to as forced parameterization.

C: Delayed transaction durability is accomplished using asynchronous log writes to disk.

Transaction log

records are kept in a buffer and written to disk when the buffer fills or a buffer flushing event takes place.

Delayed transaction durability reduces both latency and contention within the system.

Some of the cases in which you could benefit from using delayed transaction durability are:

\* You can tolerate some data loss.

\* You are experiencing a bottleneck on transaction log writes.

\* Your workloads have a high contention rate.

Reference:

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

## **NEW QUESTION: 82**

You have an Azure SQL logical server.

You run the following script.

```

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

```

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

**NOTE:** Each correct selection is worth one point

Statements	Yes	No
The <code>orders</code> table will allow only rows to be inserted.	<input type="radio"/>	<input type="radio"/>
To create additional tables in the Sales database, the <code>LEDGER = ON</code> parameter must be used.	<input type="radio"/>	<input type="radio"/>
To ensure that a timestamp is added to each row in the <code>orders</code> table, the <code>GENERATED ALWAYS</code>	<input type="radio"/>	<input type="radio"/>

**Answer:**

Statements	Yes	No
The <code>orders</code> table will allow only rows to be inserted.	<input checked="" type="radio"/>	<input type="radio"/>
To create additional tables in the Sales database, the <code>LEDGER = ON</code> parameter must be used.	<input type="radio"/>	<input checked="" type="radio"/>
To ensure that a timestamp is added to each row in the <code>orders</code> table, the <code>GENERATED ALWAYS</code>	<input type="radio"/>	<input checked="" type="radio"/>

### NEW QUESTION: 83

You have an Azure SQL database named DB 1 in the General Purpose service tier.

You need to monitor DB 1 by using SQL Insights.

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

a.


NOTE: Each correct selection is worth one point.

To collect monitoring data, use:

- A virtual machine
- An Azure function**
- The Azure Monitor agent

To store monitoring data, create:

- A Log Analytics workspace
- An Azure SQL database
- An Azure Storage account**




**Answer:**

To collect monitoring data, use:

- A virtual machine
- An Azure function**
- The Azure Monitor agent**

To store monitoring data, create:

- A Log Analytics workspace
- An Azure SQL database**
- An Azure Storage account



**NEW QUESTION: 84**

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

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

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

NOTE: Each correct selection is worth one point.

DBCC **CHECKDB** ([DB1],  with

- NOINDEX
- REPAIR\_FAST
- REPAIR\_REBUILD

- ALL\_ERRORMSGs
- NO\_INFOMSGs
- PHYSICAL\_ONLY

**Answer:**


DBCC CHECKDB ([DB1],  with

- NOINDEX
- REPAIR\_FAST
- REPAIR\_REBUILD

- ALL\_ERRORMSGs
- NO\_INFOMSGs
- PHYSICAL\_ONLY

## Explanation

Table Description automatically generated with low confidence



Reference:

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

## NEW QUESTION: 85

You have an Azure subscription.

You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance.

You need to configure the disks on the virtual machine. The solution must meet the following requirements:

- \* Minimize latency for transaction logs.
- \* Minimize the impact on IO Of the virtual machine.

Which type of disk should you use for each workload? To answer, drag the appropriate disk types to the correct workloads. Each disk type may be used once, at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



**Answer:**



Explanation

TempDB:

Transaction logs:

Graphical user interface, text, application Description automatically generated

**NEW QUESTION: 86**

You have an Azure SQL database.

You run the following PowerShell script.

```

$serverName = "SERVER1"
$resourceGroup = "RG1"
$dbName = "DB1"

Connect-AzAccount

$server = Get-AzSqlServer -ServerName $serverName -ResourceGroupName $resourceGroup

Set-AzSqlDatabaseBackupShortTermRetentionPolicy -ResourceGroupName $resourceGroup -ServerName $server -DatabaseName $dbName -RetentionDays 21

Set-AzSqlDatabaseBackupLongTermRetentionPolicy -ServerName $server -DatabaseName $dbName -ResourceGroupName $resourceGroup -WeeklyRetention P52W -YearlyRetention PSY -WeekOfYear 52

```

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

NOTE: Each correct selection is worth one point.

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

**Answer:**

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: 87**

You have a version-8.0 Azure Database for MySQL database.  
You need to identify which database queries consume the most resources.  
Which tool should you use?

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

**Answer: (SHOW ANSWER)**

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time.

Query Store simplifies performance troubleshooting by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns.

Data for all users, databases, and queries is stored in the mysql schema database in the Azure Database for MySQL instance.

Reference:

<https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store>

**NEW QUESTION: 88**

You are building an Azure virtual machine.

You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS.

You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.

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

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

## Actions

## Answer Area

a virtual disk that uses the stripe layout

a virtual disk that uses the mirror layout

a volume

a virtual disk that uses the simple layout

a storage pool

### Answer:

Answer Area
a storage pool
a virtual disk that uses the stripe layout
a vloume

1 - a storage pool

2 - a virtual disk that uses the stripe layout

3 - a vloume

Reference:

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

### NEW QUESTION: 89

You are building a database in an Azure Synapse Analytics serverless SQL pool.

You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container.

Records are structured as shown in the following sample.

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

The records contain two applicants at most.

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

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

NOTE: Each correct selection is worth one point.

▼ applications

```

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

```

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

CROSS APPLY  
OPENJSON  
OPENROWSET

Microsoft

**Answer:**

▼ 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]

```

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

CROSS APPLY  
OPENJSON  
OPENROWSET

Microsoft

GO

Reference:

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

**NEW QUESTION: 90**

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
SQL1	SQL Server on Azure Virtual Machines	<b>Not applicable</b>
db1	Microsoft SQL Server database	Hosted on SQL1
mysqlbackups	General purpose v2 storage account	<b>Not applicable</b>

You need to back up db1 to mysqlbackups, and then restore the backup to a new database named db2 that is hosted on SQL1. The solution must ensure that db1 is backed up to a stripe set.

Which three Transact-SQL statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

## Statements

## Answer Area

```
RESTORE DATABASE db2 FROM URL = URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_4.bak'  
WITH CREDENTIAL = 'sqlbackup', RECOVERY,  
    MOVE 'db1_mdf' TO  
'D:\Data\db2_mdf.mdf',  
    MOVE 'db1_log' TO  
'D:\Logs\db2_log.ldf'
```

```
BACKUP DATABASE db1  
TO URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_4.bak'  
WITH CREDENTIAL = 'sqlbackup';  
GO
```

```
RESTORE DATABASE db2 FROM URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_4.bak'  
WITH RECOVERY,  
    MOVE 'db1_mdf' TO  
'D:\Data\db2_mdf.mdf',  
    MOVE 'db1_log' TO  
'D:\Logs\db2_log.ldf'
```

```
CREATE CREDENTIAL  
[https://mysqlbackups.blob.core.windows.net  
/backups]  
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',  
SECRET = '<SAS_TOKEN>'  
GO
```

```
BACKUP DATABASE db1  
TO URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_1.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_2.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_3.bak'  
, URL =  
'https://mysqlbackups.blob.core.windows.net  
/backups/db1_4.bak'
```



```
GO
```

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY  
=  
'sqlsamplebackup'  
,SECRET = '<mystorageaccountaccesskey>';  
GO
```

**Answer:**

## Statements

```
RESTORE DATABASE db2 FROM URL = URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup', RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup';
GO
```

```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db1_log.ldf'
```

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
GO
```

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY
=
'sqlsamplebackup'
, SECRET = '<mystorageaccountaccesskey>';
GO
```

## Answer Area

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
GO
```

```
RESTORE DATABASE db2 FROM URL = URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup', RECOVERY,
MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```



## Explanation

Text Description automatically generated with low confidence

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'

GO
```

Text Description automatically generated

```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH RECOVERY,
    MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
    MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-serv>

### NEW QUESTION: 91

You have an Azure SQL managed instance.

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

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

NOTE: Each correct selection is worth one point.

RESTORE	▼	DB1	FROM	▼
DATABASE				DISK = N'\\NAS01\SQLBackups\DB1.bak';
FILE				TAPE = N'\\.\Tape0'
LOG				URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Answer:

RESTORE	▼	DB1	FROM	Microsoft ▼
DATABASE				DISK = N'\\NAS01\SQLBackups\DB1.bak';
FILE				TAPE = N'\\.\Tape0'
LOG				URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!

PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam

**questions have been updated** and **answers have been corrected** get the **newest**

PrepPdf.com DP-300 dumps with Test Engine here:

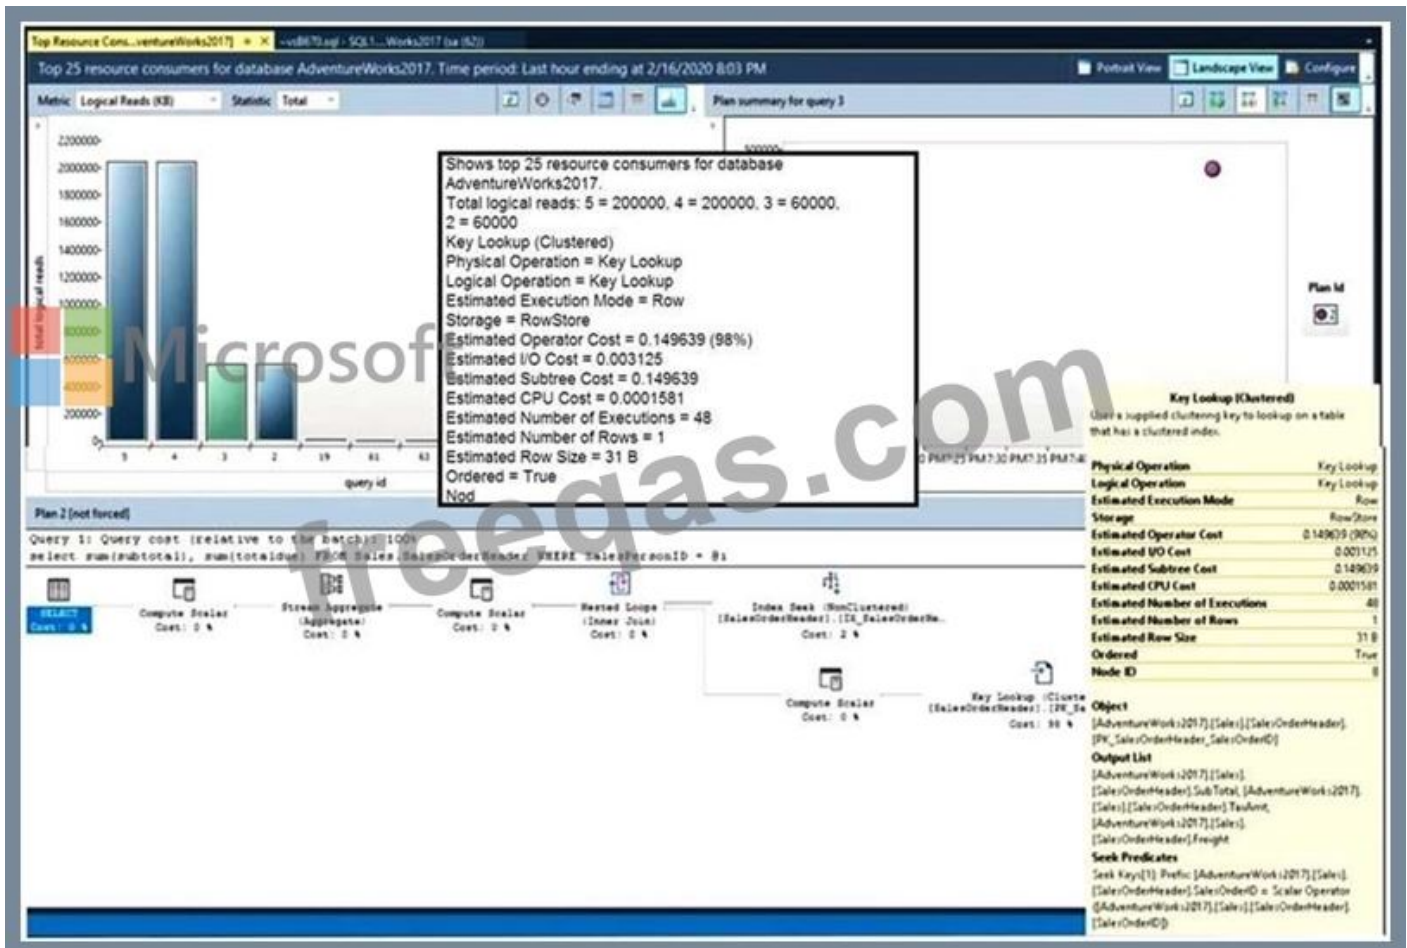
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps,

**40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 92

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



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

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<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 checked="" 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: 93**

You have a database on a SQL Server on Azure Virtual Machines instance.

The current state of Query Store for the database is shown in the following exhibit.

Microsoft

**General**

Operation Mode (Actual) Read only

Operation Mode (Requested) **Read write**

**Monitoring**

Data Flush Interval (Minutes) 15

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:**

Microsoft

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

**NEW QUESTION: 94**

You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

**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 (6 vCores)

Min vCores: 0.75, 1, 1.25, 1.5, 1.75, 2, 2.25, 2.5, 3, 4, 5, 6 (0.75 vCores)

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: 0, Hours: 4, Minutes: 0

Data max size: 1 GB, 800 GB, 1.5 TB, 800 GB (240 GB LOG SPACE ALLOCATED)

<b>Gen5 - General Purpose (GP_5, Gen5_0)</b>	
Cost per GB (in USD)	0.12
Max storage selected (in GB)	x 1040
<b>ESTIMATED STORAGE COST / MONTH</b>	<b>119.60 USD</b>
<b>COMPUTE COST / VCORE / SECOND</b>	<b>0.000145 USD</b>

**NOTES**  
1 Serverless databases are billed in vCores based on a combination of CPU and memory utilization. Learn more about serverless billing

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

NOTE: Each correct selection is worth one point.

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

**Answer:**

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

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

Reference:

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

**NEW QUESTION: 95**

You have an Azure SQL database.

You have a query and the associated execution plan as shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

The performance issue stems from the [answer choice] operator.

	▼
Select	
Index Seek	
Key Lookup	
Nested Loops	

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



	▼
heap	
clustered index	
nonclustered index	

**Answer:**

The performance issue stems from the [answer choice] operator.

	▼
Select	
Index Seek	
Key Lookup	
Nested Loops	

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

	▼
heap	
clustered index	
nonclustered index	

Explanation

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

The performance issue stems from the [answer choice] operator.

▼

- Select
- Index Seek
- Key Lookup
- Nested Loops

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

▼

- heap
- clustered index
- nonclustered index

Box 1: Key Lookup

The Key Lookup cost is 99% so that is the performance bottleneck.

Box 2: nonclustered index

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

### NEW QUESTION: 96

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

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

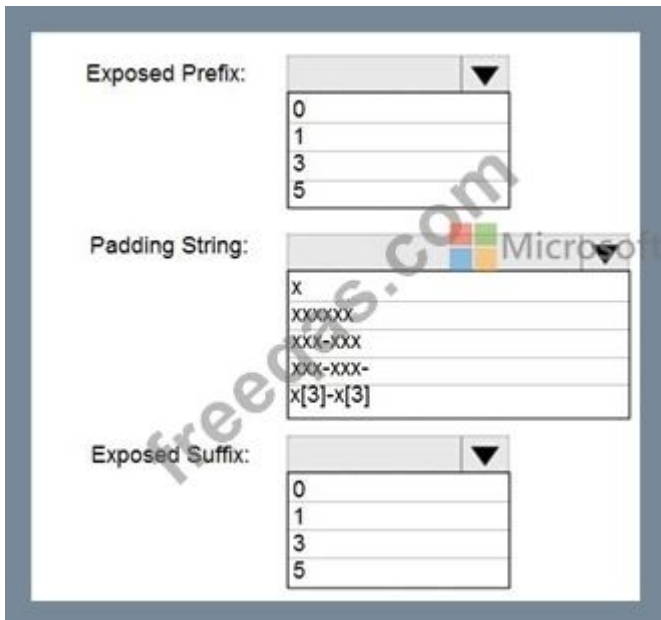
You plan to implement a dynamic data mask for the Customer\_Phone column. The mask must meet the following requirements:

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

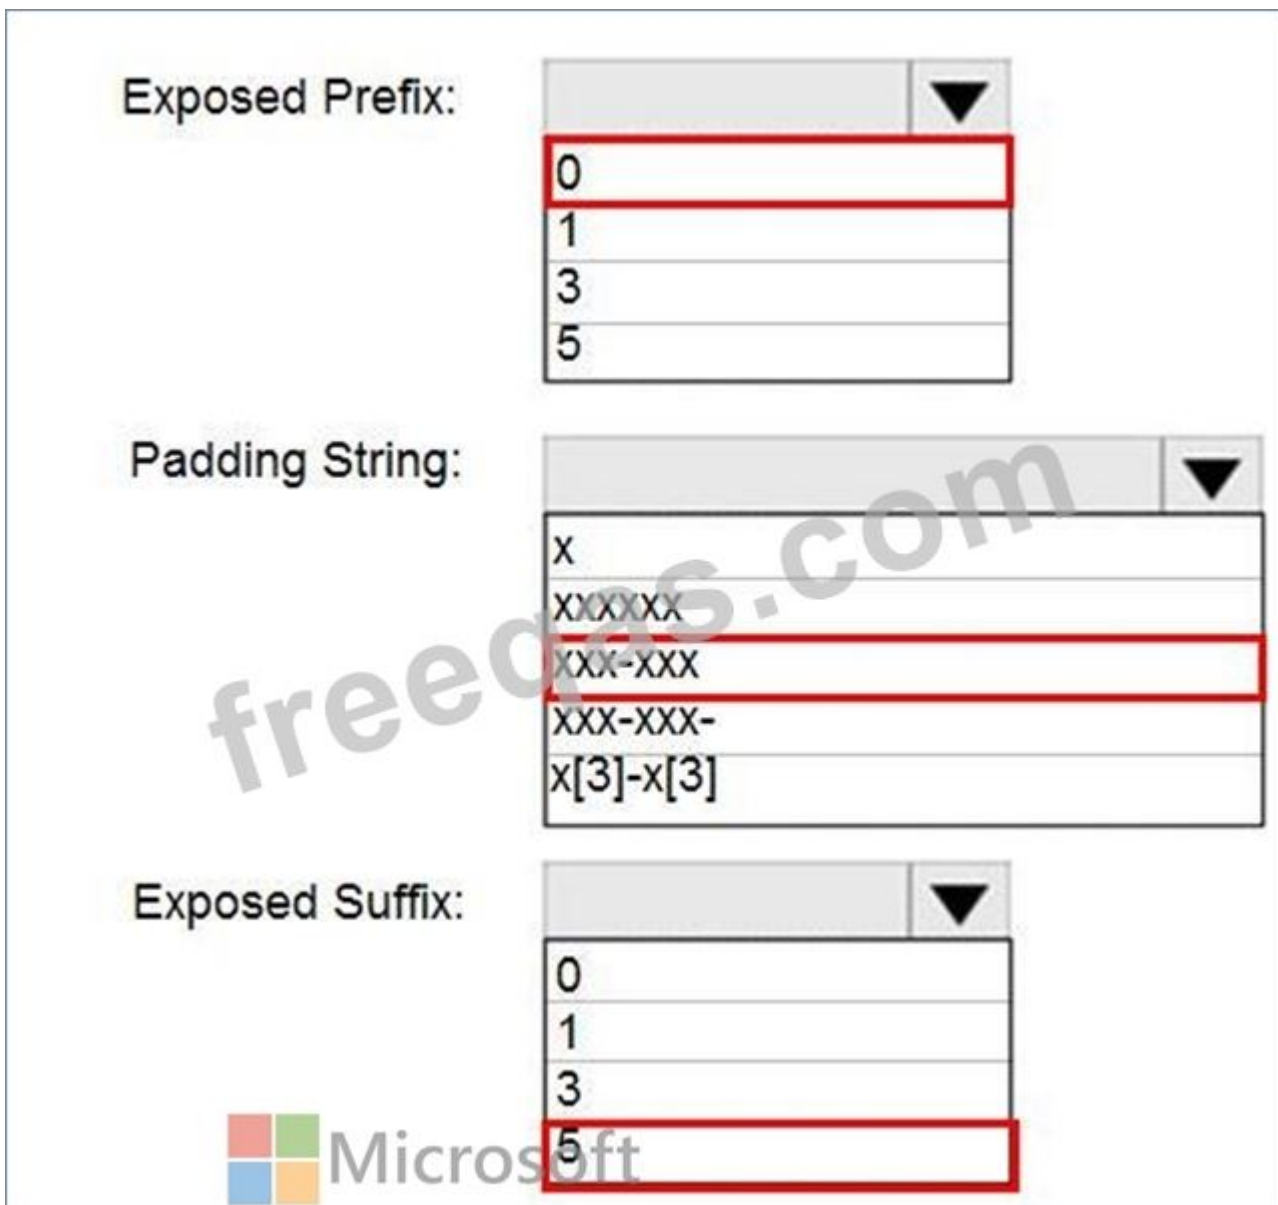
The last four digits of each customer's phone number must be visible.

Hyphens must be preserved and displayed.

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



**Answer:**



Reference:

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

**NEW QUESTION: 97**

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

You need to identify which deployment and resiliency options meet the following requirements:

Support user-initiated backups.

Support multiple automatically replicated instances across Azure regions.

Minimize administrative effort to implement and maintain business continuity.

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

NOTE: Each correct selection is worth one point.

Deployment option:

	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

Answer:

Deployment option:

	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-is-overview>

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>

**NEW QUESTION: 98**

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

- \* ProductID
- \* ItemPrice
- \* LineTotal
- \* Quantity
- \* StoreID
- \* Minute
- \* Month
- \* Hour
- \* Year
- \* Day

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
df.write
```

▼	▼
.bucketBy	("*")
.partitionBy	("StoreID", "Hour")
.range	("StoreID", "Year", "Month", "Day", "Hour")
.sortBy	("Year", "Month", "Day", "Hour", "StoreID")

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

▼
.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")

**Answer:**

# Answer Area



```
df.write
```

<pre>.bucketBy</pre>	<pre>("*")</pre>
<pre>.partitionBy</pre>	<pre>("StoreID", "Hour")</pre>
<pre>.range</pre>	<pre>("StoreID", "Year", "Month", "Day", "Hour")</pre>
<pre>.sortBy</pre>	<pre>("Year", "Month", "Day", "Hour", "StoreID")</pre>

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

<pre>.csv("/Purchases")</pre>
<pre>.json("/Purchases")</pre>
<pre>.parquet("/Purchases")</pre>
<pre>.saveAsTable("/Purchases")</pre>

Section: [none]

Explanation:

Box 1: `.partitionBy`

Example:

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

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

```
.parquet("/data/hive/warehouse/db_name.db/" + tableName)
```

Box 2: `("Year", "Month", "Day", "Hour", "StoreID")`

Box 3: `.parquet("/Purchases")`

Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partitions-with-no-new-data>

## NEW QUESTION: 99

You are building a database in an Azure Synapse Analytics serverless SQL pool.

You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container.

Records are structured as shown in the following sample.

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

The records contain two applicants at most.

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

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

NOTE: Each correct selection is worth one point.

▼ applications

CREATE EXTERNAL TABLE  
CREATE TABLE  
CREATE VIEW

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

CROSS APPLY  
OPENJSON  
OPENROWSET

Microsoft

Answer:

▼ applications Microsoft

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]
```

CROSS APPLY  
OPENJSON  
OPENROWSET

GO

Reference:

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

NEW QUESTION: 100

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

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

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

Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.

Minimize the possibility of VM1 exfiltrating data stored in SqlDb1.

D18912E1457D5D1DDCDBD40AB3BF70D5D

What should you create on VNet1?

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

**Answer: C (LEAVE A REPLY)**

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

Traffic between your virtual network and the service travels the Microsoft backbone network.

Exposing your service to the public internet is no longer necessary.

Reference:

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

### NEW QUESTION: 101

You have an Azure SQL managed instance.

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

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

NOTE: Each correct selection is worth one point.

RESTORE DATABASE DB1 FROM DISK = N'\\WAS01\\SQLBackups\\DB1.bak';  
TAPE = N'\\I\\Tape0'  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

**Answer:**

RESTORE DATABASE DB1 FROM DISK = N'\\WAS01\\SQLBackups\\DB1.bak';  
TAPE = N'\\I\\Tape0'  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

### NEW QUESTION: 102

You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1. You need to move a SQL Server Agent job from DB1 to SQLMI1.

Which job attribute is unsupported in SQLMI1?

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

Answer: ([SHOW ANSWER](#))

**NEW QUESTION: 103**

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

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

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Backup type:

Backup option:

Answer:

Backup type:

Backup option:

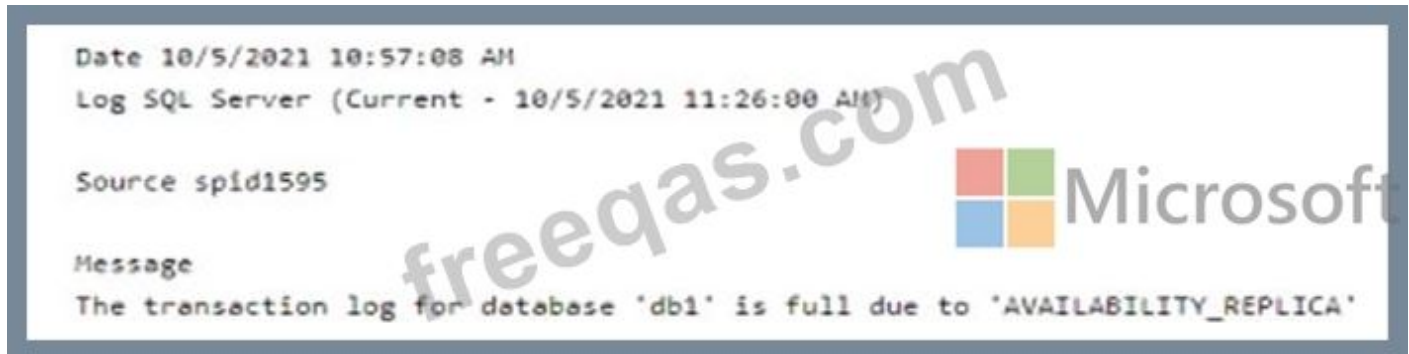
Reference:

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

**NEW QUESTION: 104**

You have a database named db1.

The log for db1 contains the following entry.



You need to ensure That db1 can process transactions.

Actions

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

Answer Area

<input type="checkbox"/>	Add db1 back to the availability group.
<input type="checkbox"/>	Shrink db1.
<input type="checkbox"/>	Shrink the transaction log file.
<input type="checkbox"/>	Remove db1 from the availability group.
<input type="checkbox"/>	Back up the transaction log file.

**Answer:**



**NEW QUESTION: 105**

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.

You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

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

**Answer: C (LEAVE A REPLY)**

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

**NEW QUESTION: 106**

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

You use Query Performance Insight to monitor db1.

You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.

Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.


NOTE: Each correct selection is worth one point.

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE

Value:

1
60
CUSTOM
ON




Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE

Value:

1
60
CUSTOM
ON



**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
 PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

**NEW QUESTION: 107**

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

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

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

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

NOTE: Each correct selection is worth one point.

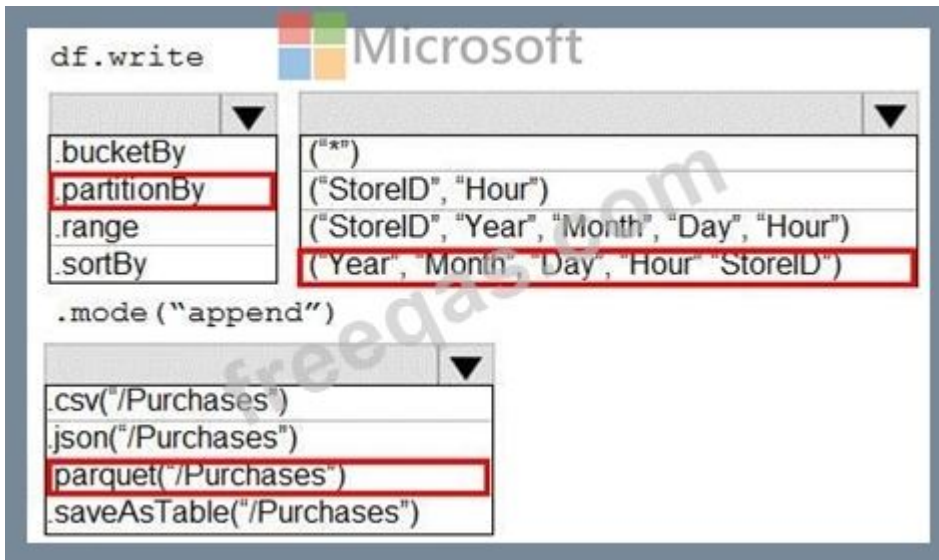
The screenshot shows a code editor with the following content:

```
df.write
```

Below the code, there are several dropdown menus for configuration options:

- bucketBy**: ("\*")
- partitionBy**: ("StoreID", "Hour")
- range**: ("StoreID", "Year", "Month", "Day", "Hour")
- sortBy**: ("Year", "Month", "Day", "Hour", "StoreID")
- mode**: ("append")
- Format**:
  - csv("/Purchases")
  - json("/Purchases")
  - parquet("/Purchases")
  - saveAsTable("/Purchases")

**Answer:**



Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partitions-with-no-new-data>


### NEW QUESTION: 108

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales.

You need to initiate a backup of Sales.

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

NOTE: Each correct selection is worth one point.



Answer:

BACKUP DATABASE Sales

TO DISK = '\\BackupSystem\BackupDisk1\Sales.bak'

TO DISK = 'X:\BAK\Sales.bak'

TO 'Sales\_Backup'

TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'

WITH STATS = 5,

WITH COPY\_ONLY;

WITH ENCRYPTION;

WITH FILE\_SNAPSHOT;

WITH NO\_TRUNCATE



Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-instance/ba-p/386154>

**NEW QUESTION: 109**

You have an Azure subscription that contains an Azure SQL database.

The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource\_semaphore waits.


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

NOTE: Each correct selection is worth one point.

```

SELECT
    is_user_process
    wait_time
    wait_type
    SUM(wait_time) AS total_wait_time_ms
FROM sys.
    dm_exec_query_stats
    dm_exec_requests
    query_store_query
JOIN sys.dm_exec_sessions AS dmvs2
    ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;

```



**Answer:**

SELECT

	▼
is_user_process	
wait_time	
wait_type	

SUM(wait\_time) AS total\_wait\_time\_ms

FROM sys.

	▼
dm_exec_query_stats	
dm_exec_requests	
query_store_query	



JOIN sys.dm\_exec\_sessions AS dmvs

ON dmvs1.session\_id = dmvs2.session\_id

WHERE is\_user\_process = 1


GROUP BY wait\_type

ORDER BY SUM(wait\_time) DESC;

Explanation

Graphical user interface, text, application Description automatically generated

```
SELECT
    is_user_process
    wait_time
    wait_type
SUM(wait_time) AS total_wait_time_ms
FROM sys.
    dm_exec_query_stats
    dm_exec_requests
    query_store_query
JOIN sys.dm_exec_sessions AS dmvs2
    ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC
```



Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

### NEW QUESTION: 110

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

You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands.

You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00.

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

### Actions

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

Wednesday differential backup

### Answer Area



### Answer:

#### Actions

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

Wednesday differential backup

### Answer Area



full backup




Wednesday differential backup

Wednesday log backups

### Explanation

full backup

Wednesday differential backup

Wednesday log backups  
 Microsoft

### NEW QUESTION: 111

You have an Azure Databricks resource.

You need to log actions that relate to changes in compute for the Databricks resource.

Which Databricks services should you log?

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

**Answer: E (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:

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

### NEW QUESTION: 112

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

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

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

NOTE: Each correct selection is worth one point.

Encryption key:  ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type:  ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

Answer:

Encryption key:  ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type:  ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

Reference:

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

**NEW QUESTION: 113**

You have an Azure SQL database.

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



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

NOTE: Each correct selection is worth one point.

The exhibit shows [answer choice]

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

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

- Index Seek
- Key Lookup
- Nested Loops

Answer:

The exhibit shows [answer choice]

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

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

- Index Seek
- Key Lookup
- Nested Loops

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15>

### NEW QUESTION: 114

You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1. You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.

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

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

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

Create a managed identity for SQLMI1.

Grant SQLMI1 read access to Azure AD.

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



Answer:

Answer Area  Microsoft

Grant SGLMI1 read access to Azure AD.

Run create Logine user1@contoso.com From External Provider on the master database.

Run tAlter server Role securityadmin Add Member user1@contoso.com

- 1 - Grant SGLMI1 read access to Azure AD.
- 2 - Run create Logine user1@contoso.com From External Provider on the master database.
- 3 - Run tAlter server Role securityadmin Add Member user1@contoso.com

Reference:

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

### NEW QUESTION: 115

You are evaluating the business goals.

Which feature should you use to provide customers with the required level of access based on their service agreement?

- A. dynamic data masking
- B. Conditional Access in Azure
- C. service principals

D. row-level security (RLS)

**Answer: D (LEAVE A REPLY)**

Reference:

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

### NEW QUESTION: 116

You have an Azure SQL managed instance.

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

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

NOTE: Each correct selection is worth one point.

RESTORE DATABASE DB1 FROM DISK = N'\\NAS01\\SQLBackups\\DB1.bak';  
TAPE = N'\\Tape0'  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

**Answer:**

RESTORE DATABASE DB1 FROM DISK = N'\\NAS01\\SQLBackups\\DB1.bak';  
TAPE = N'\\Tape0'  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

### NEW QUESTION: 117

You have an Azure subscription.

You need to deploy an Azure SQL managed instance that meets the following requirements:

\*Optimize latency.

\*Maximize the memory-to-vCore ratio.

Which service tier and hardware generation should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Service tier: Business Critical  
Business Critical  
General Purpose  
Hyperscale  
Hardware generation: Premium-series - memory optimized  
Premium-series  
Standard-series (Gen 5)  
Premium-series - memory optimized

**Answer:**

Answer Area



Service tier: Business Critical

- Business Critical
- General Purpose
- Hyperscale

Hardware generation: Premium-series - memory optimized

- Premium-series
- Standard-series (Gen 5)
- Premium-series - memory optimized

Explanation

Answer Area

Service tier: Business Critical

Hardware generation: Premium-series - memory optimized

### NEW QUESTION: 118

You have an Azure SQL database.

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

Estimated query progress: 100%

Query 1: Query cost (relative to the batch): 100%

select OrderID, PickedByPersonID, OrderDate from WideWorldImporters.Sales.Orders where CustomerID = 120

Missing Index (Impact 99.6573): CREATE NONCLUSTERED INDEX [<Name of Missing Index, sysname.>] ON [Sales].[t

**SELECT** 0.126s

**Nested Loops (Inner Join)**  
0.126s  
392 of 446 (87%)

**Index Seek (NonClustered)**  
[Orders].[FK\_Sales\_Orders\_CustomerID]  
0.126s  
392 of 446 (87%)

**Key Lookup (Clustered)**  
[Orders].[PK\_Sales\_Orders]  
0.126s  
392 of 446 (87%)

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

NOTE: Each correct selection is worth one point.

The exhibit shows [answer choice].

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

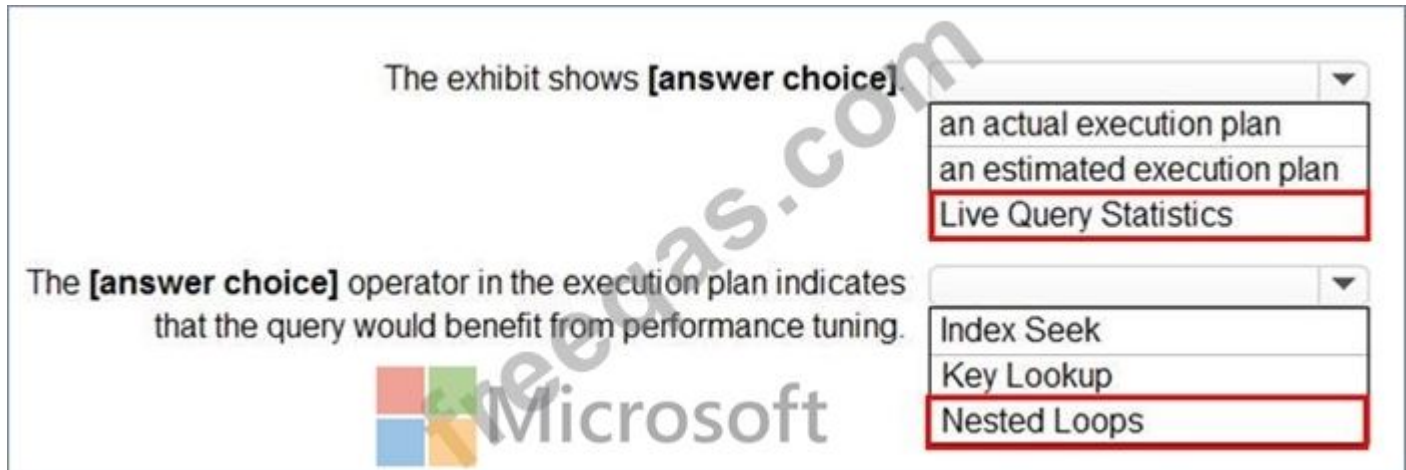
- an actual execution plan
- an estimated execution plan
- Live Query Statistics

- Index Seek
- Key Lookup
- Nested Loops

Answer:

The exhibit shows [answer choice].

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



The screenshot shows a Microsoft SQL Server interface. At the top, there is a text prompt: "The exhibit shows [answer choice]." Below this, there is a dropdown menu with three options: "an actual execution plan", "an estimated execution plan", and "Live Query Statistics". The "Live Query Statistics" option is highlighted with a red border. Below the dropdown menu, there is another text prompt: "The [answer choice] operator in the execution plan indicates that the query would benefit from performance tuning." Below this, there is another dropdown menu with three options: "Index Seek", "Key Lookup", and "Nested Loops". The "Nested Loops" option is highlighted with a red border. At the bottom of the screenshot, the Microsoft logo is visible.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15>

### NEW QUESTION: 119

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

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

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

Solution: You create additional tempdb files.

Does this meet the goal?

A. Yes

B. No

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

Reference:

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

### NEW QUESTION: 120

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

You plan to deploy an Always On availability group that will have the following configurations: VM1 will host the primary replica of DB1.

VM2 will host a secondary replica of DB1.

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

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

Microsoft SQL Server Query Editor

```
RESTORE DATABASE MyDB1  
FROM DISK = 'D:\DB1.bak'  
WITH
```

Options: NORECOVERY, RECOVERY, STANDBY

Statement type: BACKUP, CREATE, RESTORE

**Answer:**

Microsoft SQL Server Query Editor

```
RESTORE DATABASE MyDB1  
FROM DISK = 'D:\DB1.bak'  
WITH NORECOVERY  
GO
```

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server?view=sql-server-ver15>

### NEW QUESTION: 121

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Collect metrics from:

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



Stream metrics to:

▼
Azure Event Hubs
Azure Log Analytics
Azure Storage

Answer:

Collect metrics from:



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

Stream metrics to:

▼
Azure Event Hubs
Azure Log Analytics
Azure Storage

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!

PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps,

**40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 122

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

You need to ensure that users from each company can view only the data of their respective company.

Which two objects should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

**Answer:** ([SHOW ANSWER](#))

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

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

Reference:

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

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

**NEW QUESTION: 123**

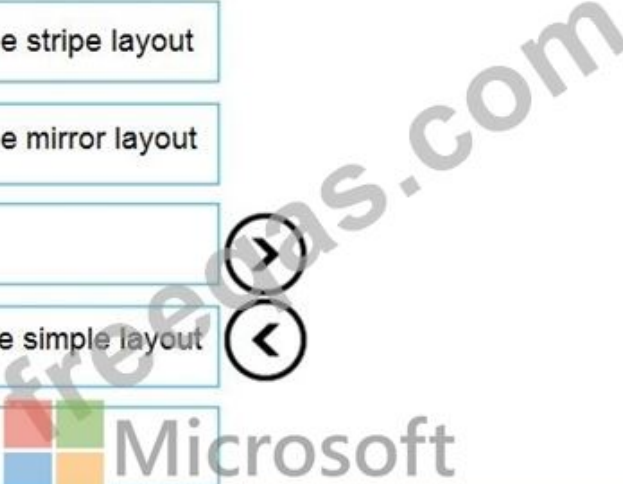

You are building an Azure virtual machine.

You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS.

You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.

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

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

Actions	Answer Area
a virtual disk that uses the stripe layout	<div style="text-align: center;">    </div>
a virtual disk that uses the mirror layout	
a volume	
a virtual disk that uses the simple layout	
a storage pool	

➤
➤
⬅
⬅
⬆
⬆

**Answer:**

**Actions**

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

**Answer Area**

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

Explanation:

Follow these same steps to create striped virtual disk:

Create Log Storage Pool.

Create Virtual Disk

Create Volume

Box 1: a storage pool

Box 2: a virtual disk that uses stripe layout

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

Box 3: a volume

Reference:

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

### NEW QUESTION: 124

You have an Azure SQL managed instance.

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

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

NOTE: Each correct selection is worth one point.

RESTORE  DATABASE  FILE  LOG DB1 FROM  DISK = N'\\NAS01\SQLBackups\DB1.bak';  TAPE = N'\\Tape0'  URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Answer:

RESTORE  DATABASE  FILE  LOG DB1 FROM  DISK = N'\\NAS01\SQLBackups\DB1.bak';  TAPE = N'\\Tape0'  URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

### NEW QUESTION: 125

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






30 TB and has a 1-GB daily rate of change.

You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands.

You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00.

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

Actions	Answer Area
Monday, Tuesday, and then Wednesday differential backups	
Wednesday, Tuesday, and then Monday log backups	
 full backup	
Monday, Tuesday, and then Wednesday log backups	
Wednesday, Tuesday, and then Monday differential backups	
Wednesday log backups	 
Wednesday differential backup	

**Answer:**

**Actions**

- Monday, Tuesday, and then Wednesday differential backups
- Wednesday, Tuesday, and then Monday log backups
- full backup
- Monday, Tuesday, and then Wednesday log backups
- Wednesday, Tuesday, and then Monday differential backups
- Wednesday log backups
- Wednesday differential backup

**Answer Area**

- full backup
- Wednesday differential backup
- Wednesday log backups

Explanation

- full backup
- Wednesday differential backup
- Wednesday log backups

**NEW QUESTION: 126**

You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

Cost summary	
Gen5 - General Purpose (GP_S, 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	0.000145 USD

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

NOTE: Each correct selection is worth one point.

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

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

**Answer:**

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

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

Explanation

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

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/database/serverless-tier-overview>

### NEW QUESTION: 127

You have an Azure virtual machine based on a custom image named VM1.

VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements:

Automate the patching of SQL Server and Windows Server.

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

Minimize administrative effort.

What should you do first?

- A. Enable a system-assigned managed identity for VM1
- B. Register VM1 to the Microsoft.Sql resource provider
- C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1  
D18912E1457D5D1DDCBD40AB3BF70D5D
- D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

**Answer: B (LEAVE A REPLY)**

Automated Patching depends on the SQL Server infrastructure as a service (IaaS) Agent Extension. The SQL Server IaaS Agent Extension (SqlIaaSExtension) runs on Azure virtual machines to automate administration tasks. The SQL Server IaaS extension is installed when you register your SQL Server VM with the SQL Server VM resource provider.

Reference:

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

### NEW QUESTION: 128

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails.

The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Premium
- B. Azure SQL Database Hyperscale

C. Azure SQL Database managed instance Business Critical

D. Azure SQL Database Standard

**Answer: (SHOW ANSWER)**

Auto-failover groups is required. SQL Managed Instance supports Auto-failover groups.

Note: The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region.

Incorrect Answers:

B: Hyperscale is for large databases. It is designed for most business workloads, providing highly scalable storage, read scale-out, and fast database restore capabilities.

Reference:

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

### NEW QUESTION: 129

You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

The screenshot shows the 'Configure' page for a new Azure SQL database. The hardware configuration is set to Gen5, with a maximum of 6 vCores and a minimum of 0.75 vCores. The data max size is set to 800 GB. The cost summary indicates a cost of 0.12 USD per GB per month for Gen5, with 1040 GB of storage selected, resulting in an estimated storage cost of 119.60 USD per month and a compute cost of 0.000145 USD per vCore per second. The auto-pause delay is currently disabled.

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

NOTE: Each correct selection is worth one point.

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

**Answer:**

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/database/serverless-tier-overview>

**NEW QUESTION: 130**

You have an Azure SQL database.

You run the following PowerShell script.

```
$serverName = "SERVER1"
$resourceGroup = "RG1"
$dbName = "DB1"

Connect-AzAccount

$server = Get-AzSqlServer -ServerName $serverName -ResourceGroupName $resourceGroup

Set-AzSqlDatabaseBackupShortTermRetentionPolicy -ResourceGroupName $resourceGroup -ServerName $server `
    -DatabaseName $dbName -RetentionDays 21

Set-AzSqlDatabaseBackupLongTermRetentionPolicy -ServerName $serverName -DatabaseName $dbName `
    -ResourceGroupName $resourceGroup -WeeklyRetention P52W -YearlyRetention PSY -WeekOfYear 52
```

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

NOTE: Each correct selection is worth one point.

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

**Answer:**

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

Text Description automatically generated

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=azsql-ps>

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

**NEW QUESTION: 131**

You have an Azure Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

You need to recommend a solution for account1 that meets the following requirements:

Automatically deletes the logs at the end of each retention period

Minimizes storage costs

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

answer area.

NOTE: Each correct selection is worth one point.

To minimize storage costs:

- Store the infrastructure logs and the application logs in the Archive access tier.
- Store the infrastructure logs and the application logs in the Cool access tier.
- Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

- Azure Data Factory pipelines
- Azure Blob storage lifecycle management rules
- Immutable Azure Blob storage time-based retention policies

**Answer:**

To minimize storage costs:

- Store the infrastructure logs and the application logs in the Archive access tier.
- Store the infrastructure logs and the application logs in the Cool access tier.
- Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

- Azure Data Factory pipelines
- Azure Blob storage lifecycle management rules
- Immutable Azure Blob storage time-based retention policies

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

### NEW QUESTION: 132

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

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

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

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

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

Hyphens must be preserved and displayed.

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

Exposed Prefix:  ▼

0
1
3
5

Padding String:  ▼

x
xxxxxx
xxx-xxx
xxx-xxx-
x[3]-x[3]

Exposed Suffix:  ▼

0
1
3
5

**Answer:**

Exposed Prefix:  ▼

0
1
3
5

Padding String:  ▼

x
xxxxxx
xxx-xxx
xxx-xxx-
x[3]-x[3]

Exposed Suffix:  ▼

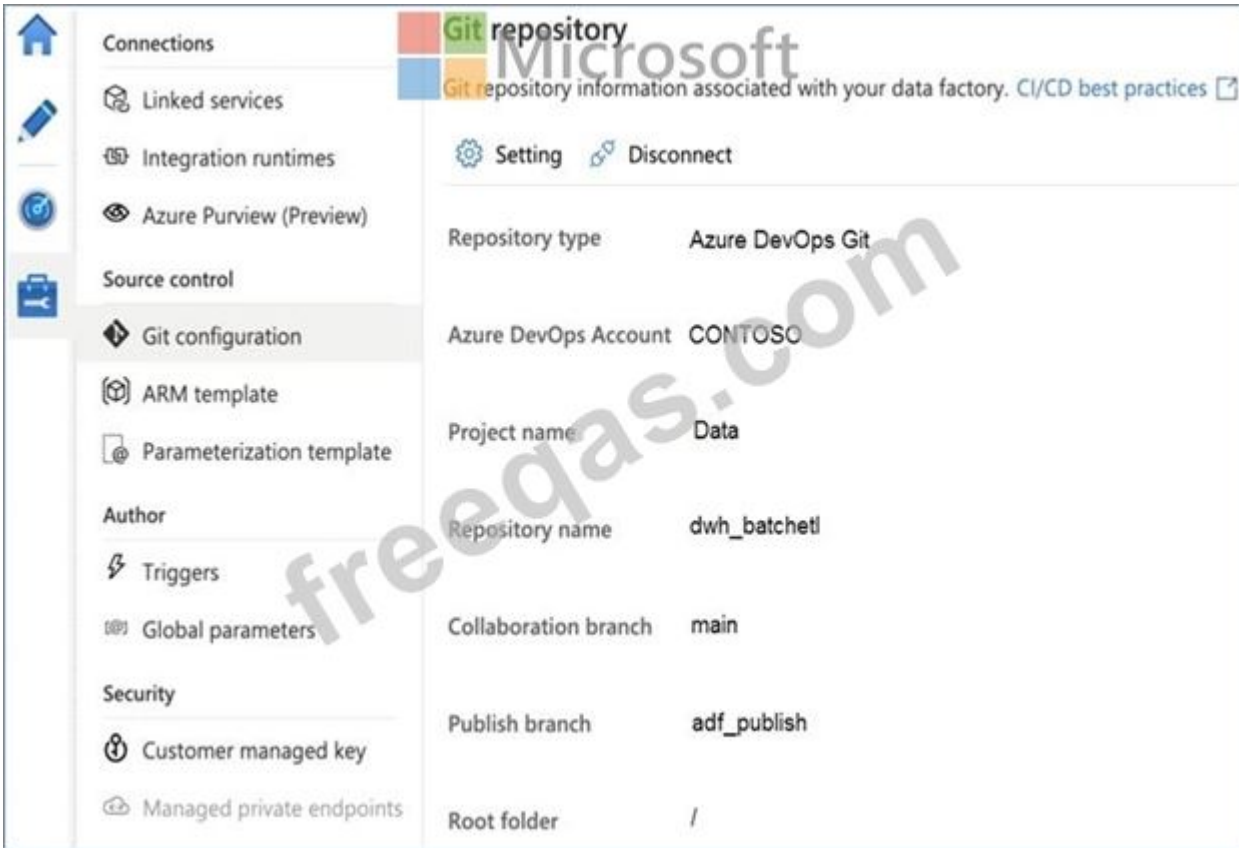
0
1
3
5

Reference:

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

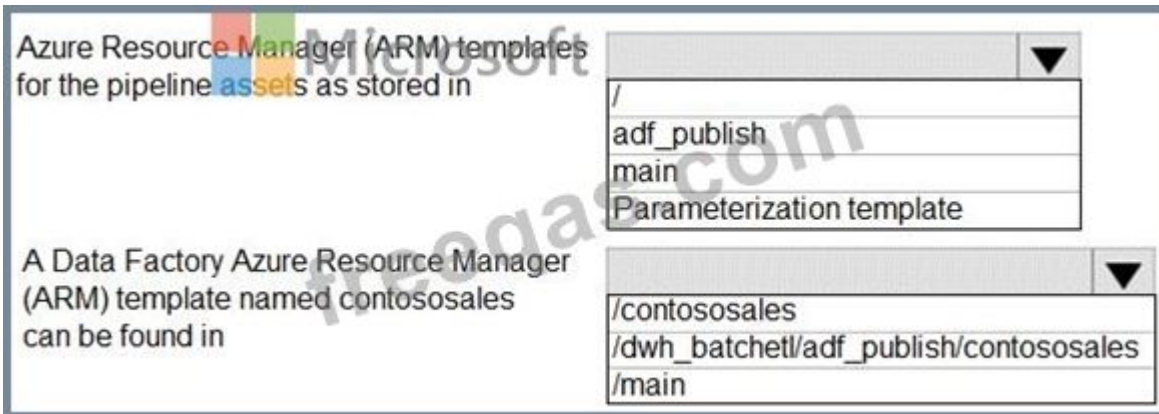
**NEW QUESTION: 133**

You configure version control for an Azure Data Factory instance as shown in the following exhibit.

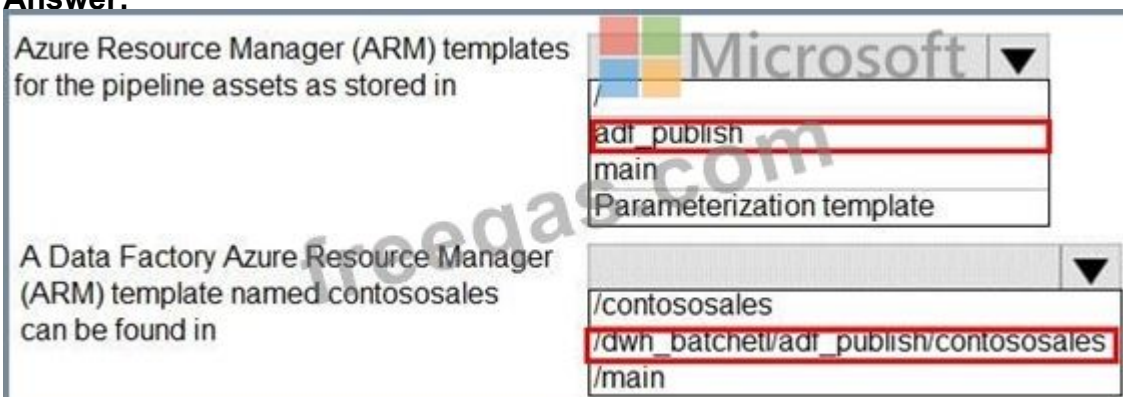


Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.



**Answer:**



Reference:

**NEW QUESTION: 134**

What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?

- A. Update the users in SalesSQLDb1.
- B. Configure SalesSQLDb1 as writable.
- C. Update the connection strings of SalesSQLDb1App1.
- D. Update the firewall rules of SalesSQLDb1.

**Answer: D (LEAVE A REPLY)**

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

Topic 3, ADatum Corporation

Existing Environment

ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.

SALESDB collects data from the stores and the website.

DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.

REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.

Requirements

Planned Changes

ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:

Migrate SALESDB and REPORTINGDB to an Azure SQL database.

Migrate DOCDB to Azure Cosmos DB.

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.

As they arrive, all the sales documents in JSON format must be transformed into one consistent format.

Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.

Technical Requirements

The new Azure data infrastructure must meet the following technical requirements:

Data in SALESDB must be encrypted by using Transparent Data Encryption (TDE). The encryption must use your own key.

SALESDB must be restorable to any given minute within the past three weeks.

Real-time processing must be monitored to ensure that workloads are sized properly based on

actual usage patterns.

Missing indexes must be created automatically for REPORTINGDB.

Disk IO, CPU, and memory usage must be monitored for SALESDB.

### NEW QUESTION: 135

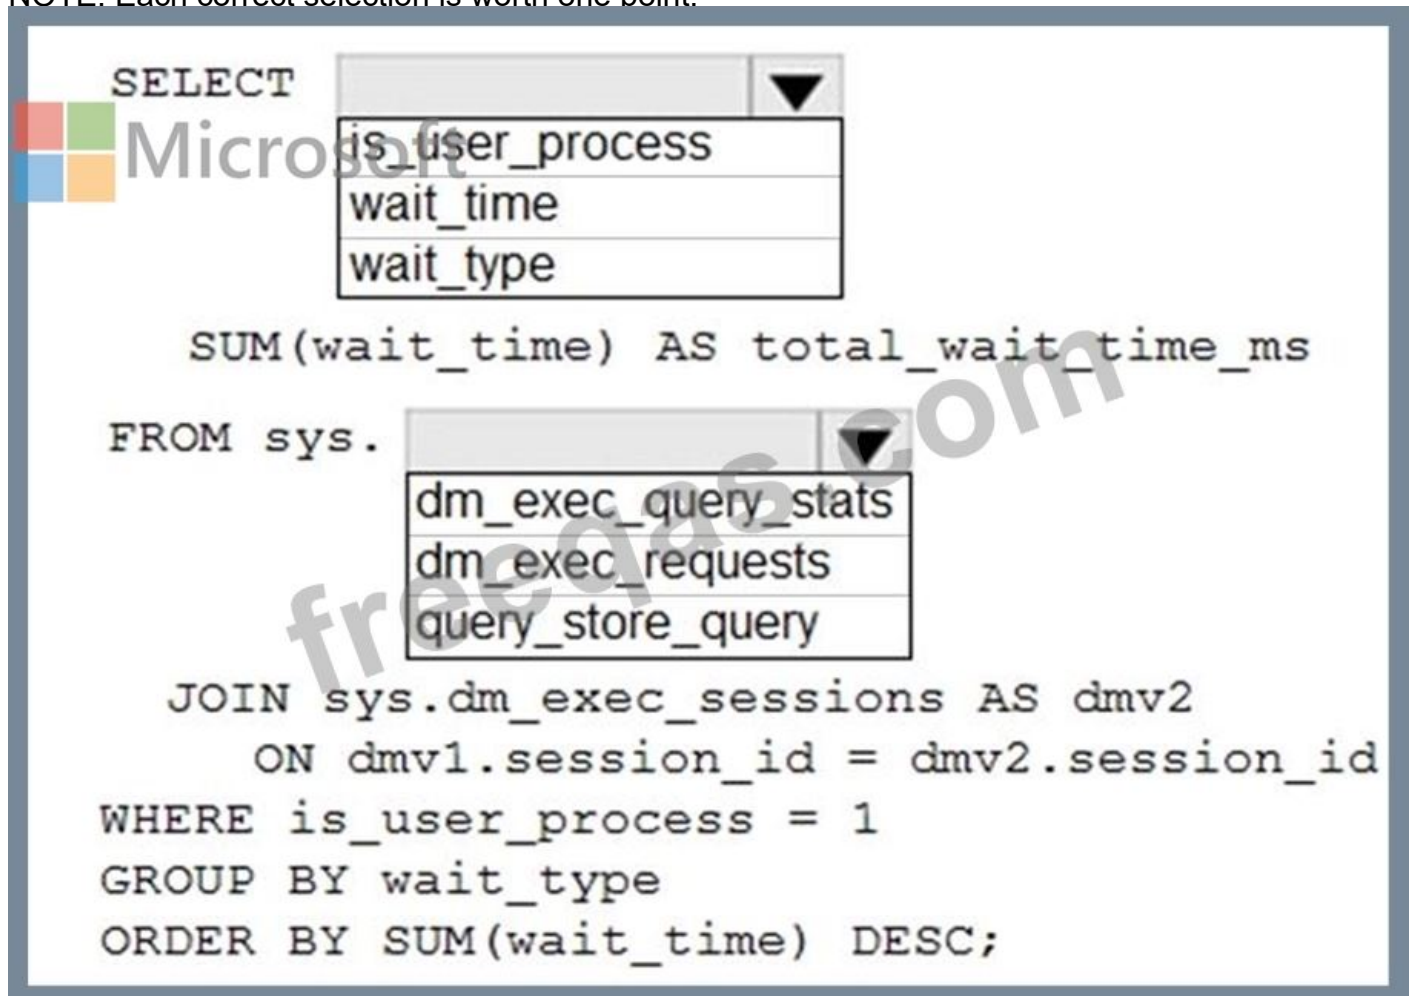
You have an Azure subscription that contains an Azure SQL database.

The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource\_semaphore waits.

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

NOTE: Each correct selection is worth one point.



The screenshot shows a SQL query editor with the following text and dropdown menus:

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

The first dropdown menu (for the SELECT clause) contains the following options: is\_user\_process, wait\_time, and wait\_type. The second dropdown menu (for the FROM clause) contains the following options: dm\_exec\_query\_stats, dm\_exec\_requests, and query\_store\_query.

**Answer:**

SELECT

is\_user\_process  
wait\_time  
wait\_type

SUM(wait\_time) AS total\_wait\_time\_ms

FROM sys.

dm\_exec\_query\_stats  
dm\_exec\_requests  
query\_store\_query

JOIN sys.dm\_exec\_sessions AS dmvs2

ON dmvs1.session\_id = dmvs2.session\_id

WHERE is\_user\_process = 1

GROUP BY wait\_type

ORDER BY SUM(wait\_time) DESC;

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

### NEW QUESTION: 136

You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named 081. You need to ensure that Group 1 has read access to new tables created in 081. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

GRANT  SELECT   [Schema]  [::table] TO [Contoso\group1]

DELETE

INSERT

SELECT

UPDATE

Database

Schema

Table

Answer:

Answer Area

GRANT  SELECT   [Schema]  [::table] TO [Contoso\group1]

DELETE

INSERT

SELECT

UPDATE

Database

Schema

Table

## Explanation

Answer Area

```
GRANT SELECT ON Schema TO [Contoso\group1]
```

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 137

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

The data contains the following columns:

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

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.

To which table should you add each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

EventCategory:  ▼

DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping:  ▼

DimChannel
DimDate
DimEvent
FactEvents

TotalEvents:  ▼

DimChannel
DimDate
DimEvent
FactEvents

**Answer:**

EventCategory:  ▼

DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping:  ▼

DimChannel
DimDate
DimEvent
FactEvents

TotalEvents:  ▼

DimChannel
DimDate
DimEvent
FactEvents

## Explanation

Graphical user interface, application, table Description automatically generated



Box 1: FactEvents

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

Box 2: DimChannel

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

Box 3: DimEvent

Reference:

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

### NEW QUESTION: 138

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

CREATE USER [app1] FROM EXTERNAL PROVIDER

CREDENTIAL FOR LOGIN app1

LOGIN FROM EXTERNAL PROVIDER

USER FROM LOGIN app1

WITHOUT LOGIN

Answer:

Answer Area

CREATE USER [app1] FROM EXTERNAL PROVIDER

CREDENTIAL FOR LOGIN app1

LOGIN FROM EXTERNAL PROVIDER

USER FROM LOGIN app1

WITHOUT LOGIN

Explanation

Answer Area

CREATE USER [app1] FROM EXTERNAL PROVIDER

**NEW QUESTION: 139**

You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.

You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.

Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

MFA: Microsoft

- Azure AD authentication
- Microsoft SQL Server authentication
- Passwordless authentication
- Windows authentication

Database-level authentication:

- Application roles
- Contained database users
- Database roles
- Microsoft SQL Server logins

Answer:



Reference:

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

#### NEW QUESTION: 140

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

You need to identify which deployment and resiliency options meet the following requirements:

Support user-initiated backups.

Support multiple automatically replicated instances across Azure regions.

Minimize administrative effort to implement and maintain business continuity.

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

NOTE: Each correct selection is worth one point.

Deployment option:

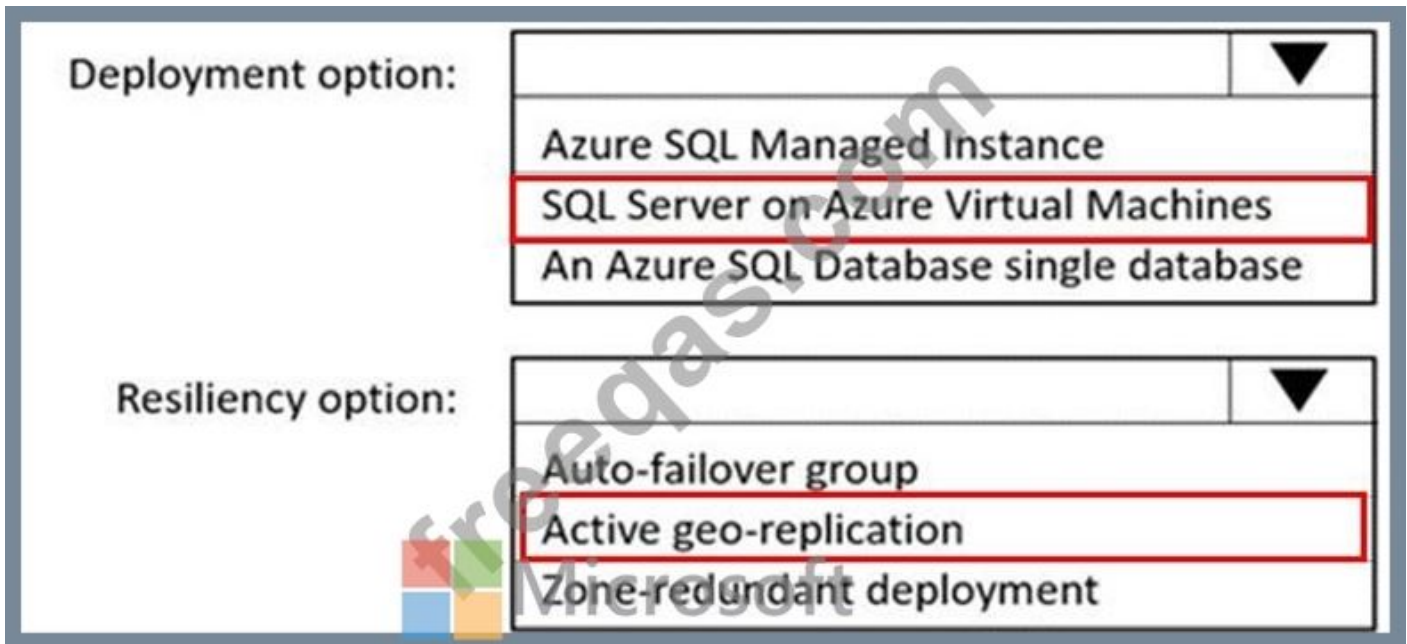
	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

A watermark 'freedqas.com' is visible across the form.

Answer:



Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-is-overview>

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>

### NEW QUESTION: 141

You have an Azure SQL Database instance named DatabaseA on a server named Server1. You plan to add a new user named App1 to DatabaseA and grant App1 db\_datacenter permissions. App1 will use SQL Server Authentication.

You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.

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

#### Actions

#### Answer Area

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

On DatabaseA, run `CREATE USER [APP1] WITH PASSWORD = 'P@ssw0rd!';`

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

On the master database, run `CREATE LOGIN [App1] WITH PASSWORD = 'P@ssw0rd!';`

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



**Answer:**

### Answer Area

On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'.

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

On DatabaseA run ALTER ROLE db\_datareader ADD Member [App1]

1 - On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'.

2 - On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]

3 - On DatabaseA run ALTER ROLE db\_datareader ADD Member [App1]

Reference:

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

### NEW QUESTION: 142

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.

Which Azure Storage functionality should you include in the solution?

- A. time-based retention
- B. change feed
- C. lifecycle management
- D. soft delete

**Answer: C (LEAVE A REPLY)**

The lifecycle management policy lets you:

Delete blobs, blob versions, and blob snapshots at the end of their lifecycles Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts>

### NEW QUESTION: 143

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

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:**

**Answer Area**

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

Section: [none]

Explanation:

Azure Database Migration service

Box 1: Premium 4-VCORE

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

\* Minimize downtime during the migration of the SERVER1 databases.

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

Incorrect Answers:

The Standard pricing tier suits most small- to medium- business workloads, but it supports offline migration only.

Box 2: A VPN gateway

You need to create a Microsoft Azure Virtual Network for the Azure Database Migration Service by using the Azure Resource Manager deployment model, which provides site-to-site connectivity to your on-premises source servers by using either ExpressRoute or VPN.

Reference:

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

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

## Testlet 1

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam.

You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

### Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

### Physical Locations

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

### Existing Environment

#### Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

#### Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com.

All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

#### Database Environment

The sales department has the following database workload:

\* An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

\* A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

\* An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1 Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### Current Problems

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

#### Requirements

#### Planned Changes

Litware plans to implement the following changes:

\* Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

\* Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

\* Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

\* Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

\* Migrate the SERVER1 databases to the Azure SQL Database platform.

#### Technical Requirements

Litware identifies the following technical requirements:

\* Maintenance tasks must be automated.

\* The 30 new databases must scale automatically.

\* The use of an on-premises infrastructure must be minimized.

\* Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

\* All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

#### Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

\* Store encryption keys in Azure Key Vault.

\* Retain backups of the PII data for two months.

\* Encrypt the PII data at rest, in transit, and in use.

- \* Use the principle of least privilege whenever possible.
- \* Authenticate database users by using Active Directory credentials.
- \* Protect Azure SQL Database instances by using database-level firewall rules.
- \* Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

**Business Requirements**

Litware identifies the following business requirements:

- \* Meet an SLA of 99.99% availability for all Azure deployments.
- \* Minimize downtime during the migration of the SERVER1 databases.
- \* Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- \* Once all requirements are met, minimize costs whenever possible.

**NEW QUESTION: 144**

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases.

The solution must meet the technical requirements and the business requirements.

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

NOTE: Each correct selection is worth one point.


Purchasing model:  ▼

- Azure virtual machine reserved instances
- DTU
- vCore

Deployment option:  ▼

- An Azure SQL Database elastic pool
- An Azure SQL Database managed instance
- A SQL Server Always On availability group

**Answer:**

Purchasing model:   ▼

- Azure virtual machine reserved instances
- DTU
- vCore

Deployment option:  ▼

- An Azure SQL Database elastic pool
- An Azure SQL Database managed instance
- A SQL Server Always On availability group

Explanation

Purchasing model:

	▼
Azure virtual machine reserved instances	
DTU	
vCore	

Deployment option:

	▼
An Azure SQL Database elastic pool	
An Azure SQL Database managed instance	
A SQL Server Always On availability group	



Microsoft

Box 1: DTU

Scenario:

The 30 new databases must scale automatically.

Once all requirements are met, minimize costs whenever possible.

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

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

Box 2: An Azure SQL database elastic pool

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

Reference:

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

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

### NEW QUESTION: 145

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails.

The solution must minimize costs.

Which deployment option and pricing tier should you configure?

A. Azure SQL Database Premium

B. Azure SQL Database serverless

C. Azure SQL Database managed instance Business Critical

D. Azure SQL Database Standard

E. This feature is not available in SQL Managed Instance.

**Answer:** ([SHOW ANSWER](#))

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

Reference:

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

Topic 1, Contoso Ltd

To start the case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

Existing Environment

Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.

Active Directory

Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises

Active Directory.

## Database Environment

Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

Name	Role
SQL1	Primary data warehouse
SQL2	Secondary data warehouse
SQL3	Extract, transform, and load (ETL) server

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.

The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.

## Current Business Model

Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers.

The customers receive the files by using FTP.

## Requirements

### Planned Changes

Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource group that contains an Azure SQL database will be provisioned for the customer. The database will have a complete copy of the financial data. The data to which each customer will have access will depend on the service agreement tier. The customers can change tiers by changing their service agreement.

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

- 

### Business Goals

Contoso identifies the following business requirements:

- 
- 
- 

### Technical Requirements

Contoso identifies the following technical requirements:

- 


### Monitoring Requirements

Contoso identifies the following monitoring requirements:

- 

### PaaS Prototype

During prototyping of the PaaS solution in Azure, you record the compute utilization of a customer's Azure SQL database as shown in the following exhibit.



**Role Assignments**

For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

Number of role assignments for this subscription: 15 (Total: 2000)

Search by name or email:  Role: 2 selected Scope: All scopes

Group by: Role

Showing a filtered set of results. Total number of role assignments: 15

2 items (2 Groups)

Name	Type	Role	Scope
DBAGroup1	Group	Contributor	This resource
DBAGroup2	Group	SQL DB Contributor	This resource

**NEW QUESTION: 146**

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool. Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted.

You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

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

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

NOTE:Each correct selection is worth one point.

Values	Answer Area
CustomerKey	<pre> 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) ) </pre>
HASH	
ROUND_ROBIN	
REPLICATE	
OrderDateKey	
SalesOrderNumber	

Answer:

Values	Answer Area
CustomerKey	<pre> 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) ) </pre>
HASH	
ROUND_ROBIN	
REPLICATE	
OrderDateKey	
SalesOrderNumber	

Explanation

Graphical user interface, text, application Description automatically generated

```

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: 147

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

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

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

NOTE: Each correct selection is worth one point.

Encryption key:  ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type:  ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

**Answer:**

Encryption key:  ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type:  ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

Reference:

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

**NEW QUESTION: 148**

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil. You plan to perform an online migration of db1 to MM by using Azure Database Migration Service.

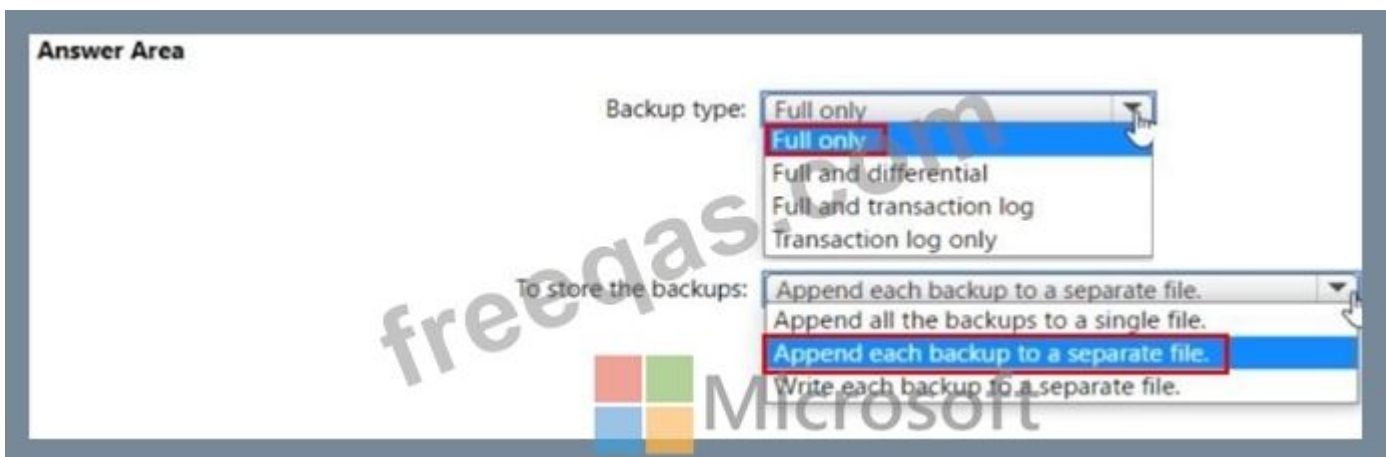
You need to create the backups for the migration. The solution must minimize the number of backup files created.

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

NOTE: Each correct selection is worth one point.



Answer:



### NEW QUESTION: 149

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

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

The records contain two applicants at most.

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

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

NOTE: Each correct selection is worth one point.

▼ applications

```

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

```

Answer:

▼ applications

```

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

```

Reference:

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

**NEW QUESTION: 150**

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

Collect metrics from:

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



Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	

Answer:

Collect metrics from:

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

Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	

**NEW QUESTION: 151**

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

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

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

Answer Area

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

Answer:

Answer Area



Statements

The log file will be truncated.

Yes

No

DB1 will be placed in an offline state.

You are performing a tail-log backup.

Explanation

Answer Area

Statements

The log file will be truncated.

DB1 will be placed in an offline state.

You are performing a tail-log backup.

Yes

No

Microsoft

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

**NEW QUESTION: 152**

You have an Azure subscription.

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

The SQL managed instance must be assigned a unique identity.

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

How should you complete the template? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Microsoft

```

{
  "type": "Microsoft.Sql/managedinstances",
  "identity": {
    "type": "UserAssigned"
  },
  "dependsOn": [
    "[parameters('virtualNetworkName')]"
  ],
  "properties": {
    "administratorLogin": "[parameters('administratorLogin')]",
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
    'subnetId')]",
    "parameters('virtualNetworkName'), parameters('subnetName')]",
    "storageLocation": "S1",
    "vCores": 80, "licenseType": "BasePrice",
    "storageAccountType": "True"
  }
}

```

**Answer:**

Microsoft

```

{
  "type": "Microsoft.Sql/managedinstances",
  "identity": {
    "type": "UserAssigned"
  },
  "dependsOn": [
    "[parameters('virtualNetworkName')]"
  ],
  "properties": {
    "administratorLogin": "[parameters('administratorLogin')]",
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
    'subnetId')]",
    "parameters('virtualNetworkName'), parameters('subnetName')]",
    "storageLocation": "S1",
    "vCores": 80, "licenseType": "BasePrice",
    "storageAccountType": "True"
  }
}

```

**NEW QUESTION: 153**

You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables.

Which distribution type should you recommend to minimize data movement?

- A. HASH
- B. REPLICATE
- C. ROUND\_ROBIN

**Answer: B (LEAVE A REPLY)**

Section: [none]

Explanation:

A replicated table has a full copy of the table available on every Compute node. Queries run fast on replicated tables since joins on replicated tables don't require data movement. Replication

requires extra storage, though, and isn't practical for large tables.

Incorrect Answers:

C: A round-robin distributed table distributes table rows evenly across all distributions. The assignment of rows to distributions is random. Unlike hash-distributed tables, rows with equal values are not guaranteed to be assigned to the same distribution.

As a result, the system sometimes needs to invoke a data movement operation to better organize your data before it can resolve a query.

Reference:

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

### NEW QUESTION: 154

You have an Azure subscription.

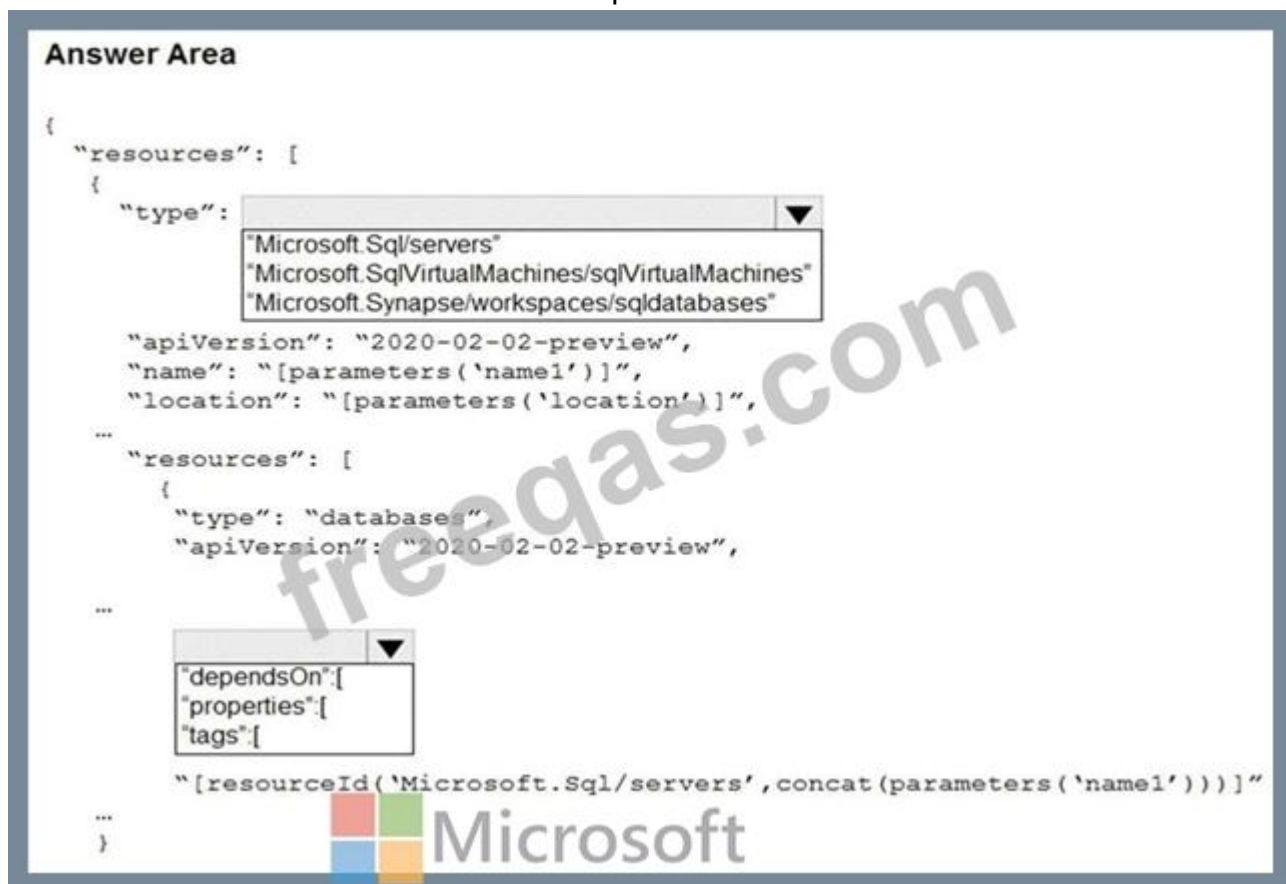
You plan to deploy an Azure SQL database by using an Azure Resource Manager template.

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

NOTE: Each correct selection is worth one point.

**Answer Area**


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



**Answer:**

## Answer Area

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



Explanation

Text Description automatically generated

```

{
  "type": [
    "Microsoft.Sql/servers",
    "Microsoft.SqlVirtualMachines/sqlVirtualMachines",
    "Microsoft.Synapse/workspaces/sqldatabases"
  ],
  "apiVersion": "2020-02-02-preview",
  "name": "[parameters('name1')]",
  "location": "[parameters('location')]",
  "resources": [
    {
      "type": "databases",
      "apiVersion": "2020-02-02-preview",
      ...
    }
  ]
}

```

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

**NEW QUESTION: 155**

You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

- A. automated
- B. interactive
- C. High Concurrency

**Answer: A (LEAVE A REPLY)**

Explanation  
 Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.  
 The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

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

**NEW QUESTION: 156**

You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1.

DF1 contains a linked service.

You have an Azure Key vault named vault1 that contains an encryption key named key1.

You need to encrypt df1 by using key1.

What should you do first?

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

**Answer: B (LEAVE A REPLY)**

Section: [none]

Explanation:

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

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

Incorrect Answers:

A, D: Should enable Soft Delete and Do Not Purge on Azure Key Vault.

Using customer-managed keys with Data Factory requires two properties to be set on the Key Vault, Soft Delete and Do Not Purge. These properties can be enabled using either PowerShell or Azure CLI on a new or existing key vault.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key> Testlet 1 This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam.

You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these

buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

#### Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

#### Physical Locations

Litware has a manufacturing office and a research office is separate locations near Boston. Each office has its own datacenter and internet connection.

#### Existing Environment

##### Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

##### Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com.

All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

##### Database Environment

The sales department has the following database workload:

- \* An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

- \* A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

- \* An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1 Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

##### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

##### Current Problems

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

##### Requirements

## Planned Changes

Litware plans to implement the following changes:

- \* Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- \* Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

- \* Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- \* Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- \* Migrate the SERVER1 databases to the Azure SQL Database platform.

## Technical Requirements

Litware identifies the following technical requirements:

- \* Maintenance tasks must be automated.
- \* The 30 new databases must scale automatically.
- \* The use of an on-premises infrastructure must be minimized.
- \* Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- \* All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

## Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- \* Store encryption keys in Azure Key Vault.
- \* Retain backups of the PII data for two months.
- \* Encrypt the PII data at rest, in transit, and in use.
- \* Use the principle of least privilege whenever possible.
- \* Authenticate database users by using Active Directory credentials.
- \* Protect Azure SQL Database instances by using database-level firewall rules.
- \* Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

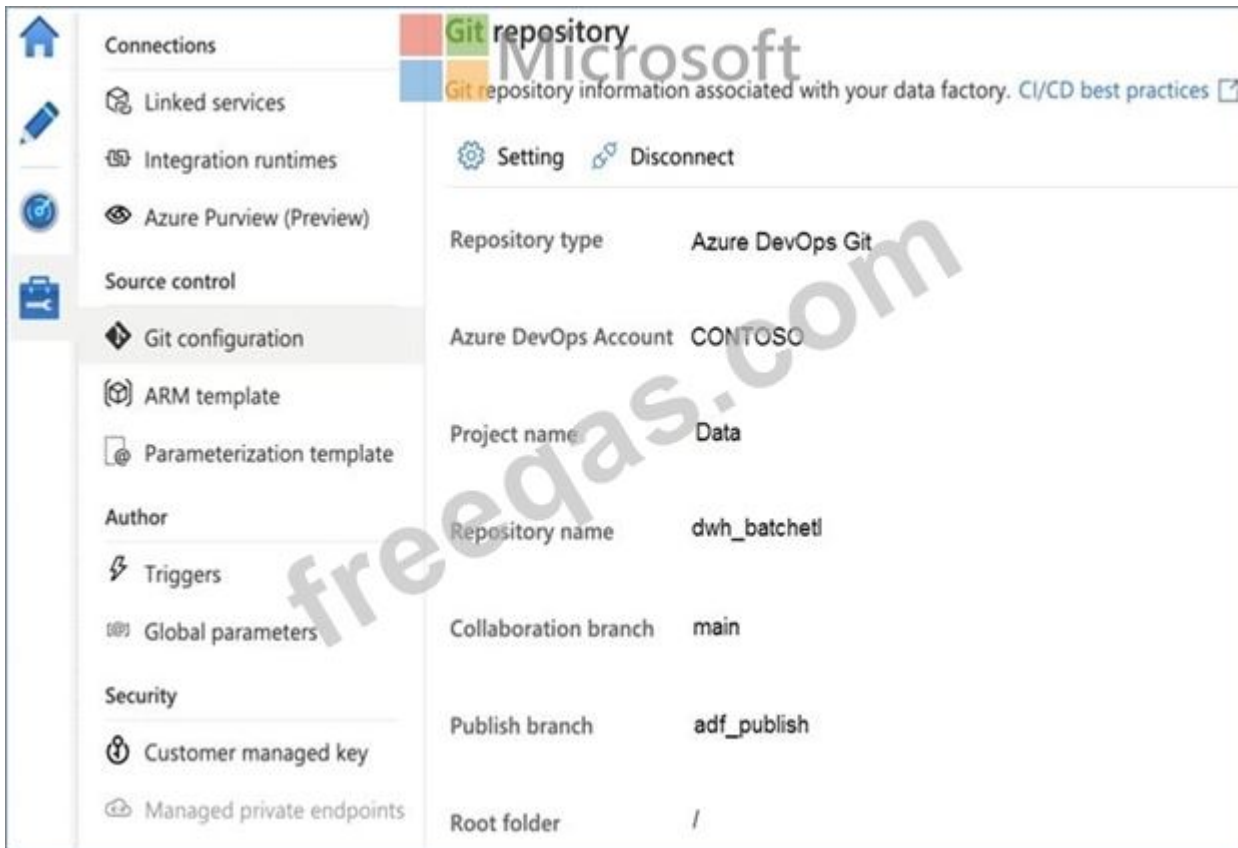
## Business Requirements

Litware identifies the following business requirements:

- \* Meet an SLA of 99.99% availability for all Azure deployments.
- \* Minimize downtime during the migration of the SERVER1 databases.
- \* Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- \* Once all requirements are met, minimize costs whenever possible.

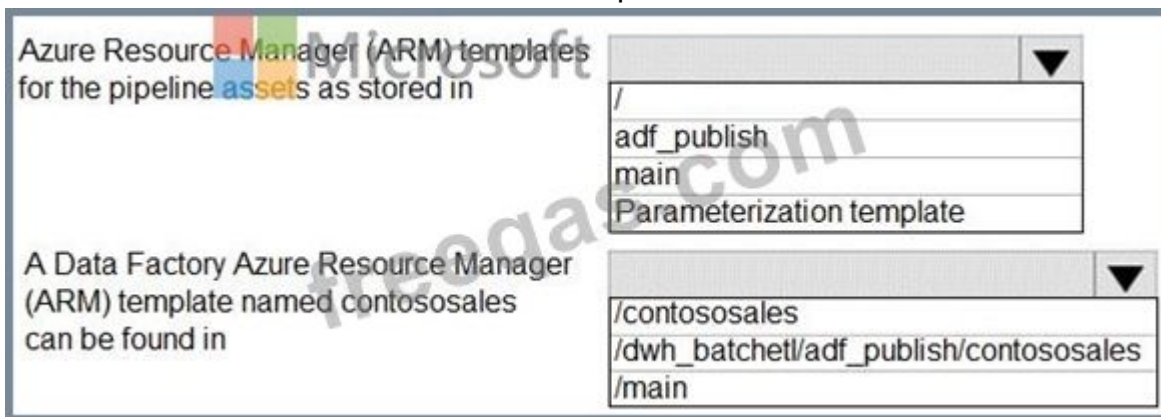
## **NEW QUESTION: 157**

You configure version control for an Azure Data Factory instance as shown in the following exhibit.

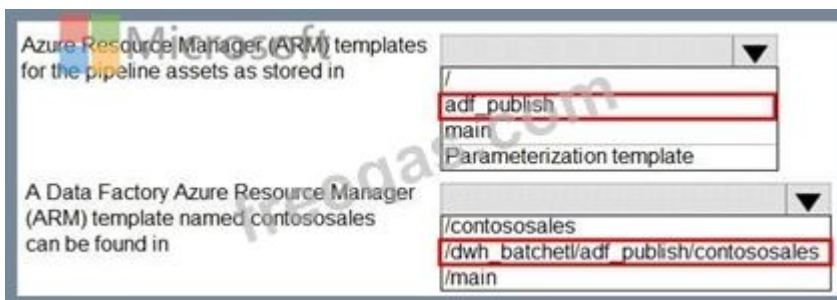


Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.



**Answer:**



Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/source-control>

**NEW QUESTION: 158**

You have an Azure SQL database that contains a table named Customer. Customer has the

columns shown in the following table.

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

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

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

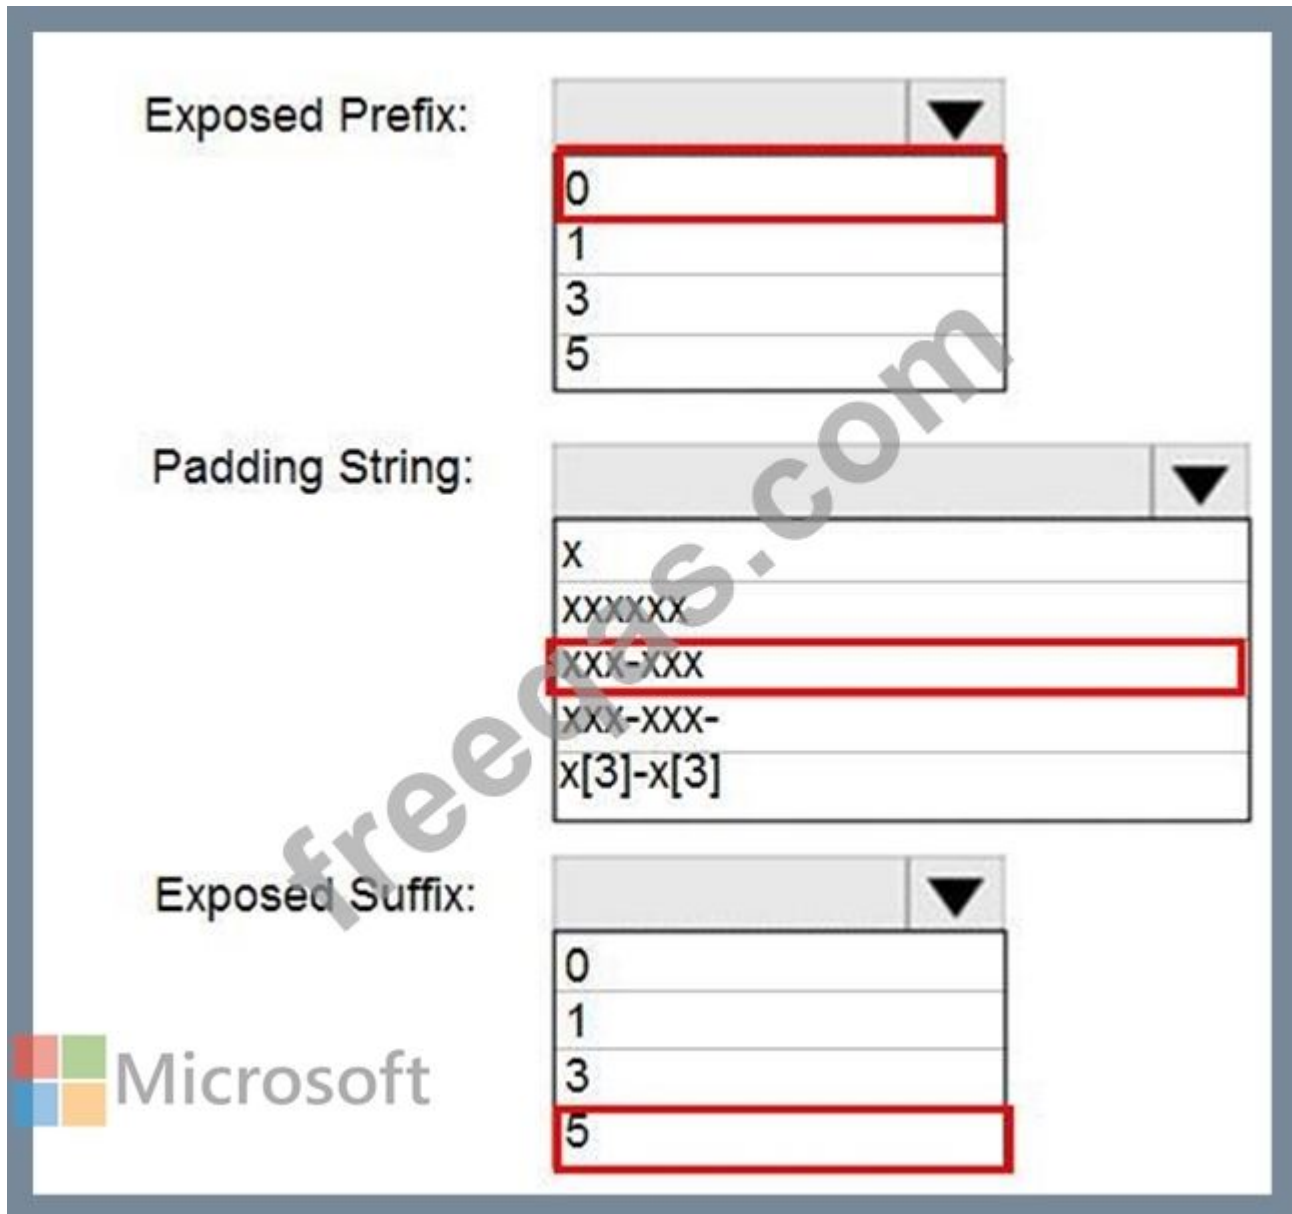
The last four digits of each customer's phone number must be visible.

Hyphens must be preserved and displayed.

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

The screenshot shows the configuration interface for a dynamic data mask. It consists of three sections:

- Exposed Prefix:** A dropdown menu with a downward arrow, containing the options 0, 1, 3, and 5.
- Padding String:** A dropdown menu with a downward arrow, containing the options X, XXXXXX, XXX-XXX, XXX-XXX-, and x[3]-x[3].
- Exposed Suffix:** A dropdown menu with a downward arrow, containing the options 0, 1, 3, and 5.



Reference:

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

**NEW QUESTION: 159**

You have an Azure SQL database.

You identify a long running query.

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

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

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

**Answer:** ([SHOW ANSWER](#))

To include an execution plan for a query during execution

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

- open an existing query and display the estimated execution plan by clicking the Open File toolbar button and locating the existing query.
2. Enter the query for which you would like to display the actual execution plan.
  3. On the Query menu, click Include Actual Execution Plan or click the Include Actual Execution Plan toolbar button.



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

Reference:

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

### NEW QUESTION: 160

You have a burstable Azure virtual machine named VM1 that hosts an instance of Microsoft SQL Server.

You need to attach an Azure ultra disk to VM1. The solution must minimize downtime on VM1. In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

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

➤    ➤    ➤    ➤    ➤    ➤

**Answer:**

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

➤    ➤    ➤    ➤    ➤

Explanation

**Answer Area** Microsoft

- 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: 161**

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Collect metrics from:

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

Stream metrics to:

- Azure Event Hubs
- Azure Log Analytics
- Azure Storage

**Answer:**

Collect metrics from:

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

Stream metrics to:

- Azure Event Hubs
- Azure Log Analytics
- Azure Storage

### NEW QUESTION: 162

Note: This question is part of a series of questions that present the same scenario. Each question in

the series contains a unique solution that might meet the stated goals. Some question sets might have

more than one correct solution, while others might not have a correct solution.

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

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

database named Database1.

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

Server2.

Solution: You run the Remove-AzSqlDatabasePowerShell cmdlet for Database1 on Server2. You run the

Restore-AzSqlDatabasePowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

A. Yes

B. No

**Answer: B (LEAVE A REPLY)**

Section: [none]

Explanation:

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

the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents

accidentally overwriting a database with a different database. If the database specified in a RESTORE

statement already exists on the current server and the specified database family GUID differs from the

database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

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

### NEW QUESTION: 163

You have an Azure SQL managed instance.

Which command should you run? To answer, select the appropriate options in the answer area. You need to restore a database named DB1 by using Transact-SQL.

NOTE: Each correct selection is worth one point.

RESTORE  DB1 FROM

DATABASE
FILE
LOG

DISK = N'\\NAS01\SQLBackups\DB1.bak';
TAPE = N'\\Tape0'
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Answer:

RESTORE  DB1 FROM

DATABASE
FILE
LOG

DISK = N'\\NAS01\SQLBackups\DB1.bak';
TAPE = N'\\Tape0'
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

**NEW QUESTION: 164**

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

The database reports a CHECKSUM error.

You need to recover the database.

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

NOTE: Each correct selection is worth one point.

```
USE master;
ALTER DATABASE [DB1] SET  WITH ROLLBACK IMMEDIATE;
GO
DBCC CHECKDB ('DB1',  WITH NO_INFOMSGS;
GO
ALTER DATABASE [DB1] SET ;
GO
```

OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY

MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST

MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;

Answer:

```

USE master;
ALTER DATABASE [DB1] SET
GO
DBCC CHECKDB ('DB1',
GO
ALTER DATABASE [DB1] SET
GO

```

Microsoft

WITH ROLLBACK IMMEDIATE;

- OFFLINE
- ONLINE
- SINGLE\_USER**
- TRUSTWORTHY

WITH NO\_INFOMSGS;

- MOINDEX
- PHYSICAL\_ONLY
- REPAIR\_ALLOW\_DATA\_LOSS
- REPAIR\_FAST**

- MULTI\_USER;**
- ONLINE;
- OPEN;
- TRUSTWORTHY;

Reference:

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

**NEW QUESTION: 165**

You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1. You need to configure the autogrow and autoshrink settings for DB1.

Which statements should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Autogrow:

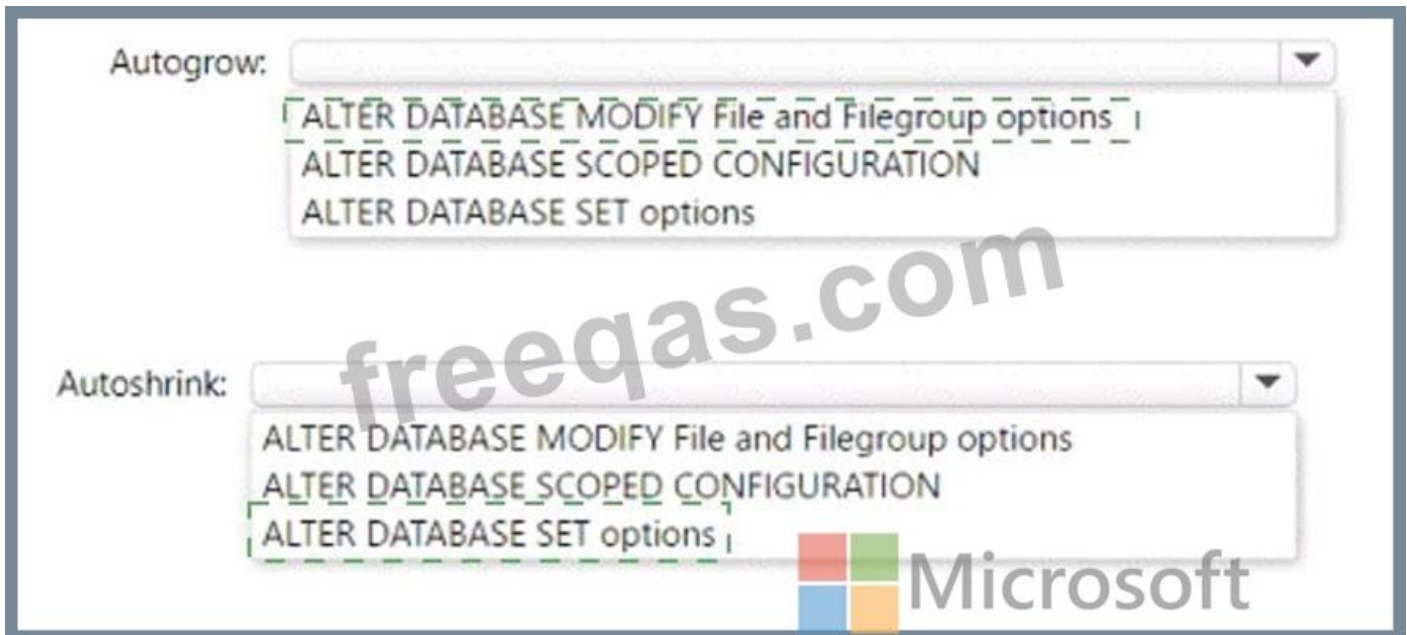
- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options



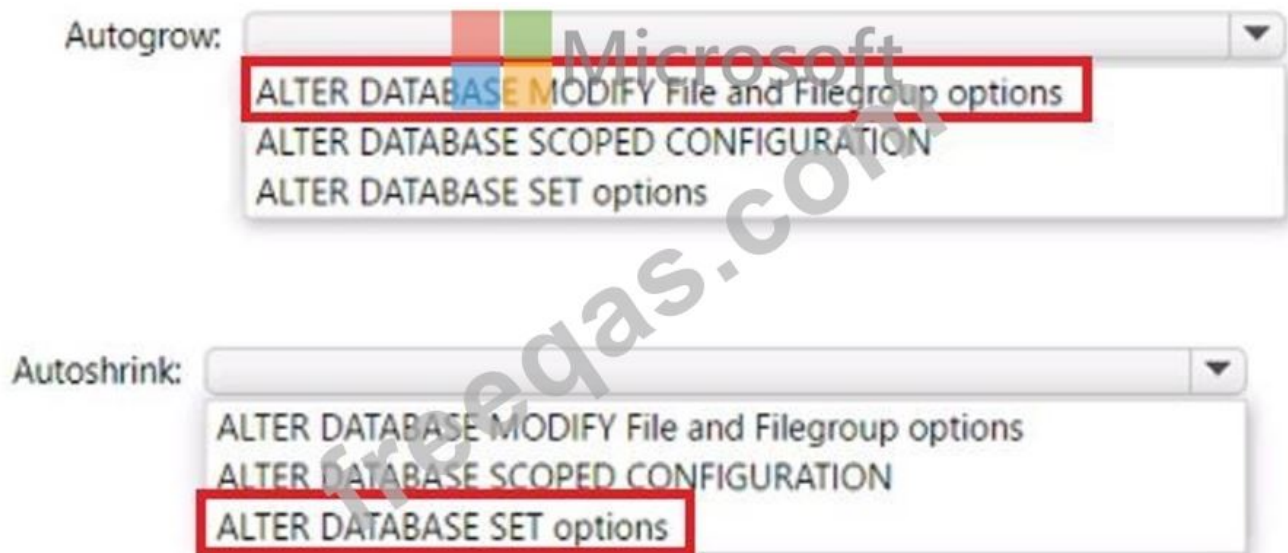
Autoshrink:

- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options

**Answer:**



Explanation

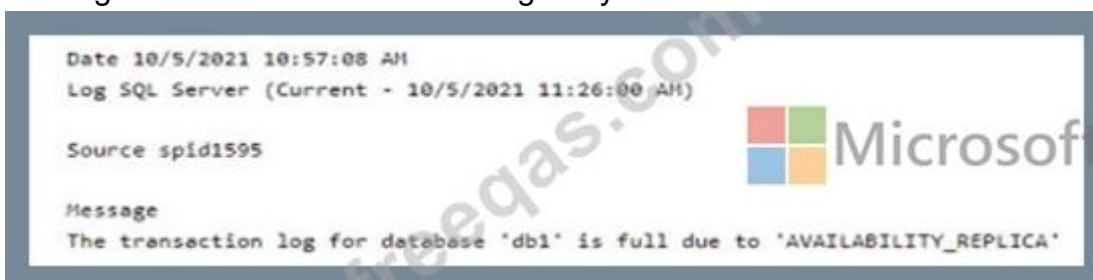


<https://learn.microsoft.com/en-us/troubleshoot/sql/admin/considerations-autogrow-autoshrink>

**NEW QUESTION: 166**

You have a database named db1.

The log for db1 contains the following entry.



You need to ensure That db1 can process transactions.

**Actions**

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

**Answer Area****Answer:**

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
 PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

**NEW QUESTION: 167**

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have SQL Server 2019 on an Azure virtual machine.

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

To gather more information, you query sys.dm\_exec\_requests and discover that the wait type is PAGELATCH\_UP and the wait\_resource is 2:3:905856.

You need to improve system performance.

Solution: You reduce the use of table variables and temporary tables.

Does this meet the goal?

**A.** Yes

**B.** No

**Answer: A (LEAVE A REPLY)**

Reference:

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

**NEW QUESTION: 168**

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine

has 4 vCPUs and 28 GB of memory.

You scale up the virtual machine to 8 vCPUSs and 64 GB of memory.

You need to provide the lowest latency for tempdb.

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

A. 2

B. 4

C. 8

D. 64

**Answer: C (LEAVE A REPLY)**

Section: [none]

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the

number of logical processors is less than or equal to eight, use the same number of data files as logical

processors. If the number of logical processors is greater than eight, use eight data files and then if contention

continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable

levels or make changes to the workload/code.

Reference:

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

Testlet 1

This is a case study. Case studies are not timed separately. You can use as much exam time as you

would like to complete each case. However, there may be additional case studies and sections on this exam.

You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the

case study. Case studies might contain exhibits and other resources that provide more information about the

scenario that is described in the case study. Each question is independent of the other questions in this case

study. At the end of this case study, a review screen will appear. This screen allows you to review your answers and to

make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore

the content of the case study before you answer the questions. Clicking these buttons displays information such

as business requirements, existing environment, and problem statements. If the case study has an All

Information tab, note that the information displayed is identical to the information displayed on the subsequent

tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com.

All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

Database Environment

The sales department has the following database workload:

- \* An on-premises server named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

- \* A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules

and contained database users.

- \* An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and

SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a

database named ManufacturingSQLDb1

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows

Server 2019 and are used to manage all the Azure databases.

### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

### Current Problems

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

### Requirements

### Planned Changes

Litware plans to implement the following changes:

- \* Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have

varying usage patterns. Each database will be approximately 20 GB.

- \* Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

- \* Develop an app named ResearchApp1 that will be used by the research department to populate and access

ResearchDB1.

- \* Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

- \* Migrate the SERVER1 databases to the Azure SQL Database platform.

### Technical Requirements

Litware identifies the following technical requirements:

- \* Maintenance tasks must be automated.

- \* The 30 new databases must scale automatically.

- \* The use of an on-premises infrastructure must be minimized.

- \* Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

- \* All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

### Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- \* Store encryption keys in Azure Key Vault.
- \* Retain backups of the PII data for two months.
- \* Encrypt the PII data at rest, in transit, and in use.
- \* Use the principle of least privilege whenever possible.
- \* Authenticate database users by using Active Directory credentials.
- \* Protect Azure SQL Database instances by using database-level firewall rules.
- \* Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

#### Business Requirements

Litware identifies the following business requirements:

- \* Meet an SLA of 99.99% availability for all Azure deployments.
- \* Minimize downtime during the migration of the SERVER1 databases.
- \* Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- \* Once all requirements are met, minimize costs whenever possible.

#### NEW QUESTION: 169

You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.

You plan to have a fact table for website visits. The table will be approximately 5 GB.


You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.

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

NOTE: Each correct selection is worth one point.

The screenshot shows two dropdown menus. The first is labeled 'Distribution:' and has a list of options: 'Hash', 'Round robin', and 'Replicated'. The second is labeled 'Index:' and has a list of options: 'Clustered columnstore', 'Clustered', and 'Nonclustered'. A watermark 'freedqas.com' is visible over the interface.

**Answer:**



Reference:

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

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

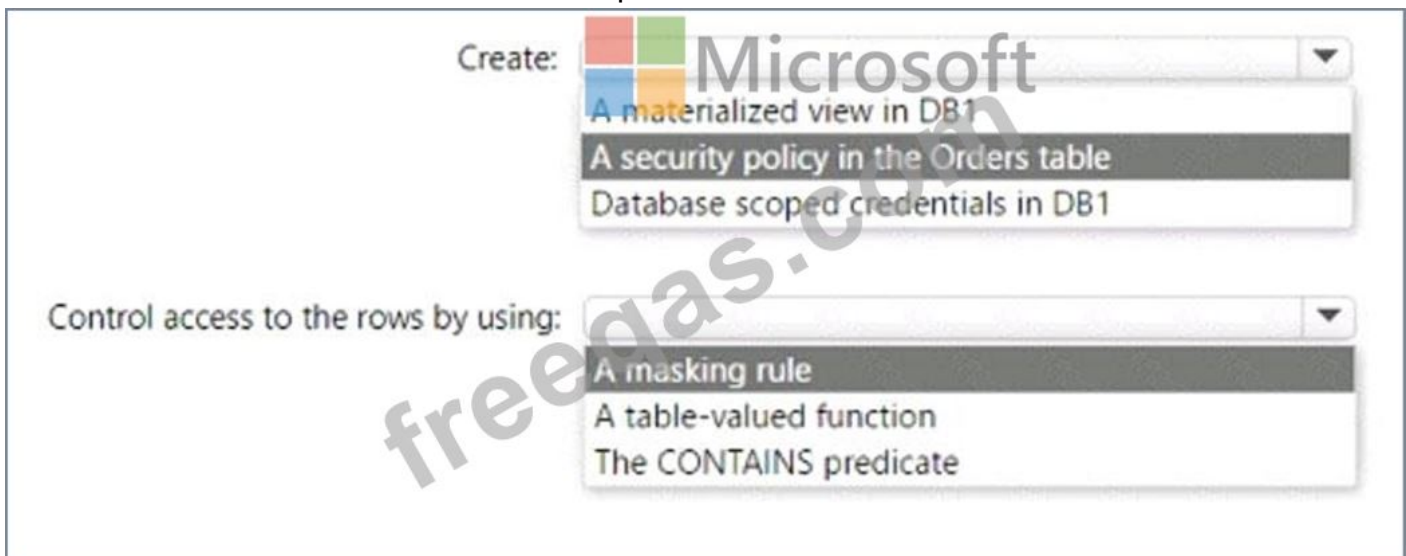
### NEW QUESTION: 170

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

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

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

NOTE: Each correct selection is worth one point.



Answer:

Create:

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

Control access to the rows by using:

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

Explanation

Create:

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

Control access to the rows by using:

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

**NEW QUESTION: 171**

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

You need to identify which deployment and resiliency options meet the following requirements:

Support user-initiated backups.

Support multiple automatically replicated instances across Azure regions.

Minimize administrative effort to implement and maintain business continuity.

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

NOTE: Each correct selection is worth one point.

Deployment option:		▼
	Azure SQL Managed Instance	
	SQL Server on Azure Virtual Machines	
	An Azure SQL Database single database	
Resiliency option:		▼
	Auto-failover group	
	Active geo-replication	
	Zone-redundant deployment	

Answer:

Deployment option:		▼
	Azure SQL Managed Instance	
	SQL Server on Azure Virtual Machines	
	An Azure SQL Database single database	
Resiliency option:		▼
	Auto-failover group	
	Active geo-replication	
	Zone-redundant deployment	

Explanation

Box 1: SQL Server on Azure VMs

SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.

Box 2: Active geo-replication

Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-is>

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>


**NEW QUESTION: 172**

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

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

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

Backup type:  Microsoft ▼

Full and log backups only
Full backup only
Log backup only

Backup option: ▼

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD

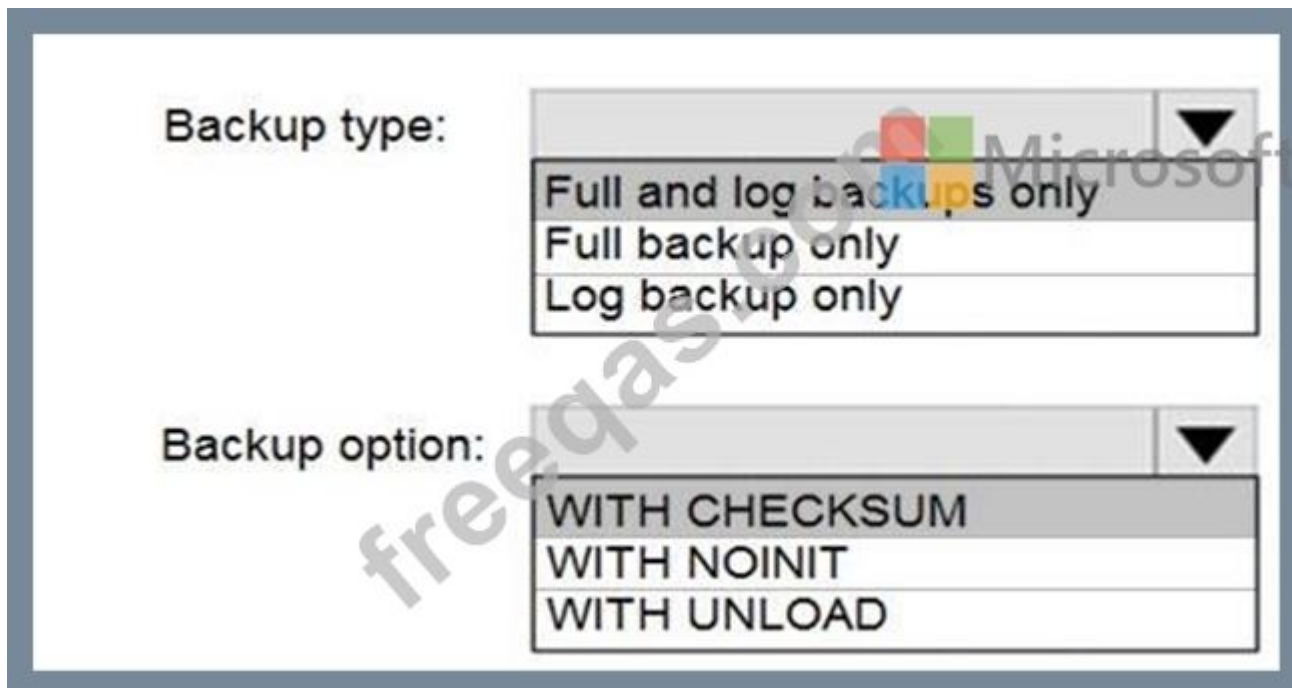
**Answer:**

Backup type: ▼

Full and log backups only
Full backup only
Log backup only

Backup option: ▼

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD



Box 1: Full and log backups only

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

Box 2: WITH CHECKSUM

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

Reference:

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

### NEW QUESTION: 173

#### HOTSPOT

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

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

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

- \* The first six numerals of each customer's phone number must be masked.
- \* The last four digits of each customer's phone number must be visible.
- \* Hyphens must be preserved and displayed.

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

Hot Area:

## Answer Area



Microsoft

Exposed Prefix:

	▼
0	
1	
3	
5	

Padding String:

	▼
x	
XXXXXX	
XXX-XXX	
XXX-XXX-	
x[3]-x[3]	

Exposed Suffix:

	▼
0	
1	
3	
5	

Answer:

## Answer Area

Exposed Prefix:

	▼
0	
1	
3	
5	

Padding String:

	▼
x	
xxxxxx	
xxx-xxx	
xxx-xxx-	
x[3]-x[3]	

Exposed Suffix:

	▼
0	
1	
3	
5	



Section: [none]

Explanation:

Box 1: 0

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

### NEW QUESTION: 174

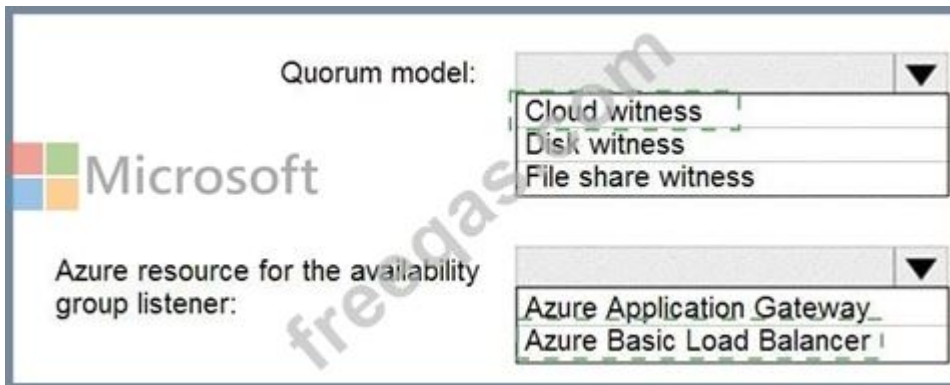
You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure. The solution must meet the business requirements.

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


NOTE: Each correct selection is worth one point.



**Answer:**



**Explanation**



**Scenario: Business Requirements**

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

**Box 1: Cloud witness**

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

**Box 2: Azure Basic Load Balancer**

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

Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a 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/](https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/)

**NEW QUESTION: 175**

You have SQL Server on an Azure virtual machine named SQL1.

SQL1 has an agent job to back up all databases.

You add a user named dbadmin1 as a SQL Server Agent operator.

You need to ensure that dbadmin1 receives an email alert if a job fails.

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

The screenshot shows a drag-and-drop interface with two columns: 'Actions' and 'Answer Area'. The 'Actions' column contains five items: 'Create a job alert', 'Create a job notification', 'Enable Database Mail', 'Enable the email settings for the SQL Server Agent', and 'Create a job target'. The 'Answer Area' is currently empty. A Microsoft logo is visible in the background. Navigation arrows are present between the columns.

**Answer:**

The screenshot shows the same drag-and-drop interface as above, but with three actions moved to the 'Answer Area' in a specific sequence. The 'Answer Area' contains three items: 'Enable the email settings for the SQL Server Agent', 'Create a job alert', and 'Create a job notification'. The 'Actions' column still contains the original five items. The 'Answer Area' items are arranged in the order they were moved, with 'Enable the email settings for the SQL Server Agent' at the top, 'Create a job alert' in the middle, and 'Create a job notification' at the bottom. Navigation arrows are present between the columns.

Explanation

Enable the email settings for the SQL Server Agent

Create a job alert



Create a job notification

Step 1: Enable the email settings for the SQL Server Agent.

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Step 2: Create a job alert

Step 3: Create a job notification

Example:

-- adds an e-mail notification for the specified alert (Test Alert)

-- This example assumes that Test Alert already exists

-- and that Francois Ajenstat is a valid operator name.

```
USE msdb ;
```

```
GO
```

```
EXEC dbo.sp_add_notification
```

```
@alert_name = N'Test Alert',
```

```
@operator_name = N'Francois Ajenstat',
```

```
@notification_method = 1 ;
```

```
GO
```

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status>

<https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator>

### **NEW QUESTION: 176**

#### **HOTSPOT**

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales.

You need to initiate a backup of Sales.

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

BACKUP DATABASE Sales

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

WITH STATS = 5,

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

Answer:

## Answer Area



BACKUP DATABASE Sales

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

WITH STATS = 5,

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

Section: [none]

Explanation:

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

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

```
BACKUP DATABASE tpcc2501
```

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

Box 2: WITH COPY\_ONLY Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-instance/ba-p/386154>

**NEW QUESTION: 177**

**HOTSPOT**

You are building an Azure Stream Analytics job to retrieve game data.

You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.

How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

SELECT  as HighestScore

- Collect(Score)
- CollectTop(1)OVER(ORDER BY Score Desc)
- Game, MAX(Score)
- TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

FROM input TIMESTAMP BY CreatedAt

GROUP BY

- Game
- Hopping(minute, 5)
- Tumbling(minute, 5)
- Windows(TumblingWindow(minute, 5), Hopping(minute, 5))

**Answer:**

**Answer Area**

SELECT  as HighestScore

- Collect(Score)
- CollectTop(1)OVER(ORDER BY Score Desc)
- Game, MAX(Score)
- TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

FROM input TIMESTAMP BY CreatedAt

GROUP BY

- Game
- Hopping(minute, 5)
- Tumbling(minute, 5)
- Windows(TumblingWindow(minute, 5), Hopping(minute, 5))

Section: [none]

Explanation:

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

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

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

<LIMIT DURATION clause> [<WHEN clause>]) Box 2: Tumbling(minute 5) Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

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



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

Reference:

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

**NEW QUESTION: 178**

HOTSPOT

You have an Azure SQL database named db1.

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

SELECT \*

FROM

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

WHERE database\_name = 'db1' AND

start\_time >

▼
DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

(day, -7, GETDATE())

ORDER BY start\_time DESC;

Answer:

Answer Area

SELECT \*

FROM

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

WHERE database\_name = 'db1' AND

start\_time >

▼
DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

(day, -7, GETDATE())

ORDER BY start\_time DESC;

Section: [none]

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
```

Incorrect Answers:

sys.dm\_exec\_requests:

sys.dm\_exec\_requests returns information about each request that is executing in SQL Server. It does not

have a column named database\_name.

sys.dm\_db\_resource\_stats:

sys.dm\_db\_resource\_stats does not have any start\_time column.

Note: sys.dm\_db\_resource\_stats returns CPU, I/O, and memory consumption for an Azure SQL Database

database. One row exists for every 15 seconds, even if there is no activity in the database.

Historical data is

maintained for approximately one hour.

Sys.dm\_user\_db\_resource\_governance returns actual configuration and capacity settings used by resource

governance mechanisms in the current database or elastic pool. It does not have any start\_time column.

Reference:

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

### **NEW QUESTION: 179**

You have an Azure SQL logical server.

You run the following script.

```

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

```



freeqas.com

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

**NOTE:** Each correct selection is worth one point

**Statements**

The `orders` table will allow only rows to be inserted.

**Yes**

**No**

To create additional tables in the Sales database, the `LEDGER = ON` parameter must be used.



To ensure that a timestamp is added to each row in the `orders` table, the `GENERATED ALWAYS`



**Answer:**

**Statements**

The `orders` table will allow only rows to be inserted.

**Yes**

**No**

To create additional tables in the Sales database, the `LEDGER = ON` parameter must be used.



To ensure that a timestamp is added to each row in the `orders` table, the `GENERATED ALWAYS`



**NEW QUESTION: 180**

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

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

**NOTE:** Each correct selection is worth one point.



**Answer:**



Reference:

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

### **NEW QUESTION: 181**

You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.

What should you include in the plan?

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

**Answer: D (LEAVE A REPLY)**

Reference:

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

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:

**NEW QUESTION: 182**

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

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

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

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

Solution: You run the Remove-AzSqlDatabasePowerShell cmdlet for Database1 on Server2. You run the Restore-AzSqlDatabasePowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

- A. Yes
- B. No

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

Section: [none]

Explanation:

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

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

Reference:

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

**NEW QUESTION: 183**

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

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

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

----- 14

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

**Weekly LTR Backups**  
Keep weekly backups for:

52 Week(s) ▾

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

60 Week(s) ▾

**Yearly LTR Backups**  
Keep an annual backup for:

10 Year(s) ▾

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

Week 52



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

NOTE: Each correct selection is worth one point.

**Answer Area**

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

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

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

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



**Answer:**

**Answer Area**


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

----- 14

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

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.

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



**NEW QUESTION: 184**

What should you implement to meet the disaster recovery requirements for the PaaS solution?

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

**Answer: (SHOW ANSWER)**

Explanation

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

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

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

Reference:

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

**NEW QUESTION: 185**

You have an Azure SQL database named DB 1 in the General Purpose service tier.

You need to monitor DB 1 by using SQL Insights.

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


NOTE: Each correct selection is worth one point.

To collect monitoring data, use:

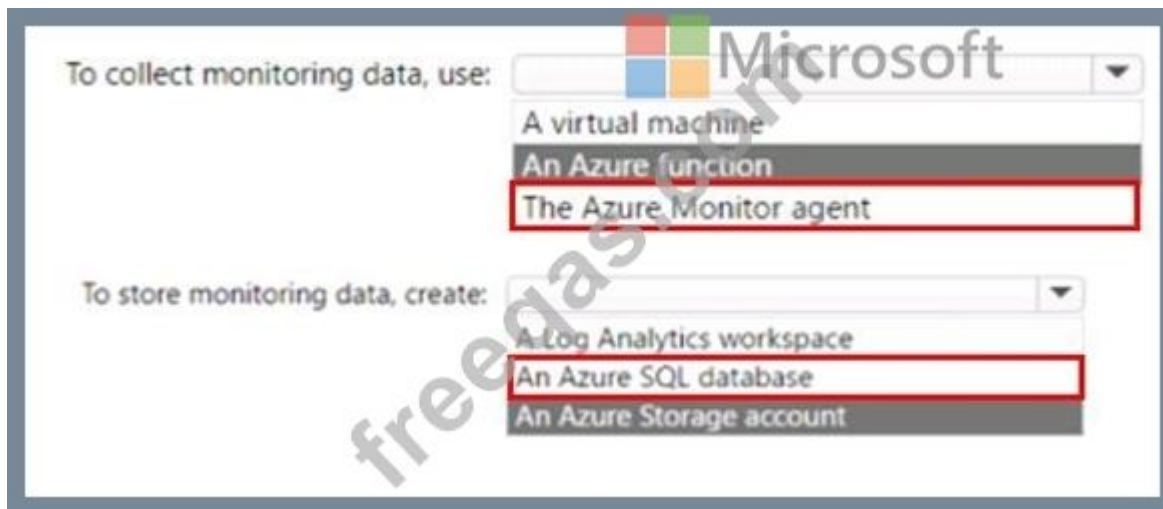
- A virtual machine
- An Azure function**
- The Azure Monitor agent

To store monitoring data, create:

- A Log Analytics workspace
- An Azure SQL database
- An Azure Storage account**



**Answer:**



To collect monitoring data, use:

- A virtual machine
- An Azure function
- The Azure Monitor agent**

To store monitoring data, create:

- A Log Analytics workspace
- An Azure SQL database**
- An Azure Storage account

**NEW QUESTION: 186**

You have SQL Server on Azure virtual machines in an availability group.

You have a database named DB1 that is

You create a full database backup of DB1.

You need to add DB1 to the availability group.

Which restore option should you use on the secondary replica?

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

**Answer: B (LEAVE A REPLY)**

Explanation

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

1. Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY
2. Join the restored database to the availability group.

Reference:

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

**NEW QUESTION: 187**

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

**Configure policies** ×

SQL server

**Point in Time Restore Configuration**

Configure PiTR backup retention ▼ Days

---

**Long-term Retention Configurations**

Weekly LTR Backups ⓘ

How long would you like weekly backups to be kept?

6 ✓ Week(s) ▼

---

Monthly LTR Backups ⓘ

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

12 ✓ Month(s) ▼

---

Yearly LTR Backups ⓘ

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

Week 2 ▼

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

10 ✓ Year(s) ▼

Microsoft

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

- January 4, 2020
- January 11, 2020
- January 18, 2020
- January 25, 2020
- February 1, 2020
- February 8, 2020

February 15, 2020  
February 22, 2020  
February 29, 2020  
March 7, 2020

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

NOTE: Each correct selection is worth one point.

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

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

Answer:

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

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

**NEW QUESTION: 188**

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1.

DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.

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

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

**ACTIONS**

**Answer Area**

Change the Active Directory Admin on TestServer1

Change the server name and related variables in the templates

From the database project, deploy the database schema and permissions

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

From the Azure portal, deploy the templates.



**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**

From the Azure portal, export the Azure Resource Manager templates
Change the server name and related variables in the templates
From the Azure portal, deploy the templates.
From the database project, deploy the database schema and permissions

**NEW QUESTION: 189**

You have SQL Server on Azure virtual machines in an availability group.

You have a database named DB1 that is NOT in the availability group.

You create a full database backup of DB1.

You need to add DB1 to the availability group.

Which restore option should you use on the secondary replica?

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

**Answer: B (LEAVE A REPLY)**

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

1. Restore a recent database backup of the primary database and subsequent log backups onto each server

instance that hosts the secondary replica, using RESTORE WITH NORECOVERY

2. Join the restored database to the availability group.

Reference:

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

## NEW QUESTION: 190

You have the following Transact-SQL query.

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

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

- A. ColumnA
- B. ColumnB
- C. ColumnC
- D. ColumnD

**Answer: C (LEAVE A REPLY)**

Section: [none]

Explanation:

Example:

Free space for the file in the below query result set will be returned by the FreeSpaceMB column.

```
SELECT DB_NAME() AS DbName,  
name AS FileName,  
type_desc,  
size/128.0 AS CurrentSizeMB,  
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB  
FROM sys.database_files  
WHERE type IN (0,1);
```

Reference:

<https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/>

### NEW QUESTION: 191

You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.

You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication.

You need to ensure that App1 can access DB2. The solution must meet the following requirements:

App1 must be able to view only DB2.

Administrative effort must be minimized.

What should you create?

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

**Answer: D (LEAVE A REPLY)**

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-database-portable?view=sql-server-ver15>

### NEW QUESTION: 192

You have an Azure Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

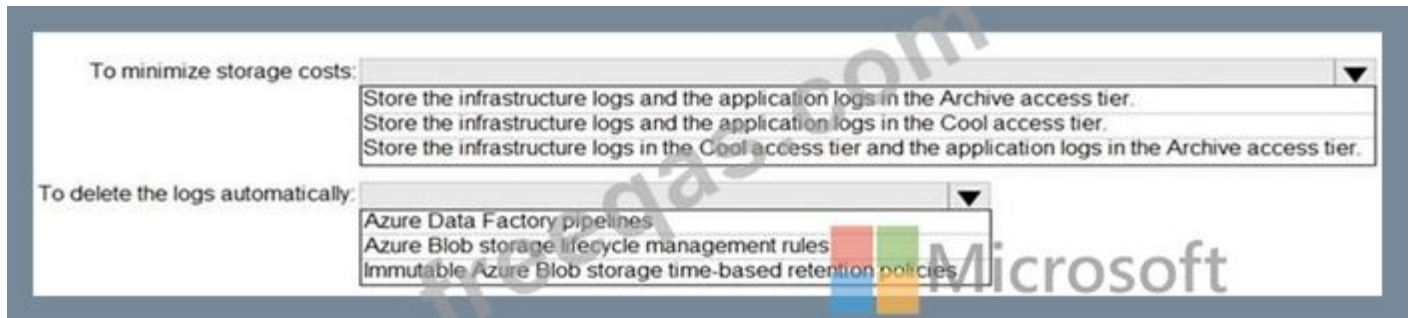
You need to recommend a solution for account1 that meets the following requirements:

Automatically deletes the logs at the end of each retention period

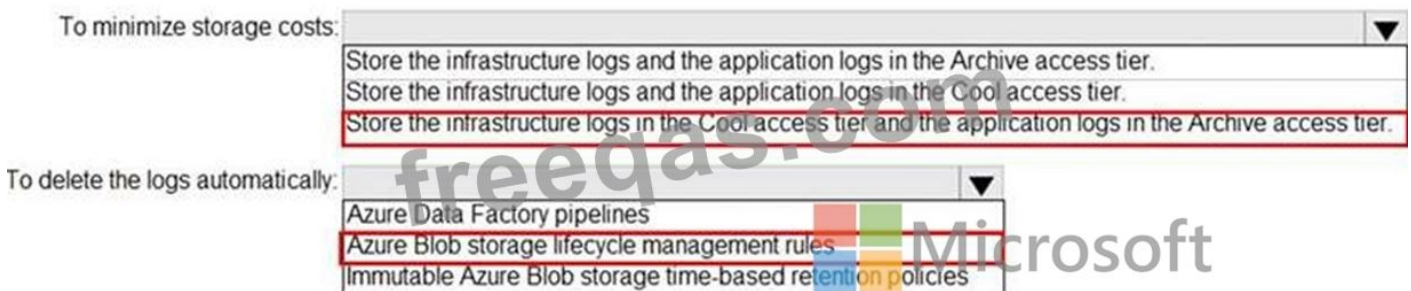
Minimizes storage costs

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

NOTE: Each correct selection is worth one point.



**Answer:**



Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

### NEW QUESTION: 193

You have an Azure SQL database.

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



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

NOTE: Each correct selection is worth one point.

The exhibit shows [answer choice]

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

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

- Index Seek
- Key Lookup
- Nested Loops

Answer:

The exhibit shows [answer choice].

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

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

- Index Seek
- Key Lookup
- Nested Loops

Explanation

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

The exhibit shows [answer choice].

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

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

- Index Seek
- Key Lookup
- Nested Loops

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver1>

### NEW QUESTION: 194

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.

You need to configure Azure Synapse Analytics to receive the data.

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

Actions	Answer Area
Create an external data source for Azure Blob storage.	
Create a master key on database.	
Enable Transparent Data Encryption.	
Create the external table FactSalesOrderDetails.	
Load the data to a staging table.	
Create an external file format to map the parquet files.	

**Answer:**

Answer Area
Create a master key on database.
Create an external data source for Azure Blob storage.
Create an external file format to map the parquet files.
Create an external table FactSalesOrderDetails

- 1 - Create a master key on database.
- 2 - Create an external data source for Azure Blob storage.
- 3 - Create an external file format to map the parquet files.
- 4 - Create an external table FactSalesOrderDetails

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure->

blob-storage

**NEW QUESTION: 195**

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

The database reports a CHECKSUM error.

You need to recover the database.

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

NOTE: Each correct selection is worth one point.

The screenshot shows a SQL query editor with three statements, each followed by a dropdown menu and a semicolon. The first statement is 'USE master;' followed by 'ALTER DATABASE [DB1] SET' and a dropdown menu with options: OFFLINE, ONLINE, SINGLE\_USER, and TRUSTWORTHY. The second statement is 'DBCC CHECKDB ('DB1',' followed by a dropdown menu with options: MOINDEX, PHYSICAL\_ONLY, REPAIR\_ALLOW\_DATA\_LOSS, and REPAIR\_FAST. The third statement is 'ALTER DATABASE [DB1] SET' followed by a dropdown menu with options: MULTI\_USER;, ONLINE;, OPEN;, and TRUSTWORTHY;. The text 'WITH ROLLBACK IMMEDIATE;' is to the right of the first dropdown, and 'WITH NO\_INFOMSGS;' is to the right of the second dropdown. A Microsoft logo is visible in the bottom right corner of the editor area.

```
USE master;  
ALTER DATABASE [DB1] SET [ ] WITH ROLLBACK IMMEDIATE;  
GO  
DBCC CHECKDB ('DB1', [ ]) WITH NO_INFOMSGS;  
GO  
ALTER DATABASE [DB1] SET [ ]  
GO
```

**Answer:**

USE master;

ALTER DATABASE [DB1] SET

GO

▼
OFFLINE
ONLINE
<b>SINGLE USER</b>
TRUSTWORTHY

WITH ROLLBACK IMMEDIATE;

DBCC CHECKDB ('DB1',

GO

▼
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
<b>REPAIR_FAST</b>

WITH NO\_INFOMSGS;

ALTER DATABASE [DB1] SET

GO

▼
<b>MULTI USER;</b>
ONLINE;
OPEN;
TRUSTWORTHY;



Microsoft

Reference:

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

### NEW QUESTION: 196

HOTSPOT

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

<b>Answer Area</b>						
Collect metrics from:	<table border="1"><tr><td>▼</td></tr><tr><td>The database only</td></tr><tr><td>The elastic pool and the database</td></tr><tr><td>The elastic pool only</td></tr><tr><td>The server, the elastic pool, and the database</td></tr></table>	▼	The database only	The elastic pool and the database	The elastic pool only	The server, the elastic pool, and the database
▼						
The database only						
The elastic pool and the database						
The elastic pool only						
The server, the elastic pool, and the database						
Stream metrics to:	<table border="1"><tr><td>▼</td></tr><tr><td>Azure Event Hubs</td></tr><tr><td>Azure Log Analytics</td></tr><tr><td>Azure Storage</td></tr></table>	▼	Azure Event Hubs	Azure Log Analytics	Azure Storage	
▼						
Azure Event Hubs						
Azure Log Analytics						
Azure Storage						

Answer:


## Answer Area

Collect metrics from:

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

Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	



Section: [none]

Explanation:

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

Scenario:

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and

storage usage and limits must be analyzed by using Azure built-in functionality.

Box 2: Azure Event hubs

Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Event hubs are able to handle custom metrics.

Incorrect Answers:

Azure Log Analytics

Azure metric and log data are sent to Azure Monitor Logs, previously known as Azure Log Analytics, directly by

Azure. Azure SQL Analytics is a cloud only monitoring solution supporting streaming of diagnostics telemetry

for all of your Azure SQL databases.

However, because Azure SQL Analytics does not use agents to connect to Azure Monitor, it does not support

monitoring of SQL Server hosted on-premises or in virtual machines.

Testlet 2

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you

would like to complete each case. However, there may be additional case studies and sections on this exam.

You must manage your time to ensure that you are able to complete all questions included on this

exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore

the content of the case study before you answer the questions. Clicking these buttons displays information such

as business requirements, existing environment, and problem statements. If the case study has an All

Information tab, note that the information displayed is identical to the information displayed on the subsequent

tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

Existing Environment

Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.

Active Directory

Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active

Directory.

Database Environment

Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

Name	Role
SQL1	Primary data warehouse
SQL2	Secondary data warehouse
SQL3	Extract, transform, and load (ETL) server

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs,

provides

historical data, and handles the delivery of data to customers.

The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.

#### Current Business Model

Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers.

The

customers receive the files by using FTP.

#### Requirements

##### Planned Changes

Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a

service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource

group that contains an Azure SQL database will be provisioned for the customer. The database will have a

complete copy of the financial data. The data to which each customer will have access will depend on the

service agreement tier. The customers can change tiers by changing their service agreement.

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

- \* Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.
- \* Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months.
- \* Start onboarding customers to the new PaaS solution within six months.

##### Business Goals

Contoso identifies the following business requirements:

- \* Use built-in Azure features whenever possible.
- \* Minimize development effort whenever possible.
- \* Minimize the compute costs of the PaaS solutions.
- \* Provide all the customers with their own copy of the database by using the PaaS solution.
- \* Provide the customers with different table and row access based on the customer's service agreement.
- \* In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.
- \* Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso.

##### Technical Requirements

Contoso identifies the following technical requirements:

- \* Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or

have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure

AD of Contoso to minimize guest users.

\* All customers must have their own resource group, Azure SQL server, and Azure SQL database. The

deployment of resources for each customer must be done in a consistent fashion.

\* Users must be able to review the queries issued against the PaaS databases and identify any new objects created.

\* Downtime during the PostgreSQL database migration must be minimized.

### Monitoring Requirements

Contoso identifies the following monitoring requirements:

\* Notify administrators when a PaaS database has a higher than average CPU usage.

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

\* Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.


\* Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues

automatically whenever possible.

### PaaS Prototype

During prototyping of the PaaS solution in Azure, you record the compute utilization of a customer's Azure SQL

database as shown in the following exhibit.



### Role Assignments

For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

Microsoft

Check access | Role assignments | Deny assignments | Classic administrators | Roles

Manage access to Azure resources for users, groups, service principals and managed identities at this scope by creating role assignments. [Learn more](#)

Number of role assignments for this subscription ⓘ

15 / 2000

Name ⓘ Type ⓘ Role ⓘ Scope ⓘ

Group by ⓘ

ⓘ Showing a filtered set of results. Total number of role assignments: 15

2 items (2 Groups)

<input type="checkbox"/>	Name	Type	Role	Scope
<input type="checkbox"/>	<b>DB</b> DBAGroup1	Group	Contributor <span>ⓘ</span>	This resource
<input type="checkbox"/>	<b>DB</b> DBAGroup2	Group	SQL DB Contributor <span>ⓘ</span>	This resource

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam! PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 197

You have an Azure SQL database.

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

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

NOTE: Each correct selection is worth one point.

Answer Area



Microsoft

```
SELECT wait_type
       wait_type
       context_info
       wait_resource
       SUM(wait_time) AS total_wait_time_ms
FROM sys.dm_exec_requests AS dev1
JOIN sys.dm_exec_sessions AS sess
     ON dev1.session_id = sess.session_id
JOIN sys.dm_db_partition_stats AS part_stats
     ON dev1.partition_id = part_stats.partition_id
WHERE is_user_process = 1
GROUP BY TARGET1
ORDER BY SUM(wait_time) DESC;
```

Answer:

Answer Area

```
SELECT wait_type
       wait_type
       context_info
       wait_resource
       SUM(wait_time) AS total_wait_time_ms
FROM sys.dm_exec_requests AS dev1
JOIN sys.dm_exec_sessions AS sess
     ON dev1.session_id = sess.session_id
JOIN sys.dm_db_partition_stats AS part_stats
     ON dev1.partition_id = part_stats.partition_id
WHERE is_user_process = 1
GROUP BY TARGET1
ORDER BY SUM(wait_time) DESC;
```

Explanation

```
SELECT wait_type
       SUM(wait_time) AS total_wait_time_ms
FROM sys.dm_exec_requests AS dev1
JOIN sys.dm_exec_sessions AS sess
     ON dev1.session_id = sess.session_id
WHERE is_user_process = 1
GROUP BY TARGET1
ORDER BY SUM(wait_time) DESC;
```

**NEW QUESTION: 198**

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

- ProductID
- ItemPrice
- LineTotal
- Quantity
- StoreID
- Minute
- Month
- Hour

Year

Day

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

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

NOTE: Each correct selection is worth one point.

```
df.write
.bucketBy
.partitionBy
.range
.sortBy
.mode("append")
.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")
```

**Answer:**

```
df.write
.bucketBy
.partitionBy
.range
.sortBy
.mode("append")
.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")
```

Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partitions-with-no-new-data>

### NEW QUESTION: 199

You have an Azure SQL database named db1.

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

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

NOTE: Each correct selection is worth one point.

SELECT \*

FROM

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

WHERE database\_name = 'db1' AND

start\_time >

▼
DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

(day, -7, GETDATE())

ORDER BY start\_time DESC;

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 >
    DATEADD
    DATEDIFF
    DATEPART
    TODATETIMEOFFSET
    (day, -7, GETDATE())
ORDER BY start_time DESC;
```

Reference:

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

### NEW QUESTION: 200

You need to use an Azure Resource Manager (ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O performance for the SQL Server database and log files. How should you complete the template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
"variables": {
  "dataDisks": {
    "caching": [
      None
      ReadOnly
      ReadWrite
    ]
  }
}

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

Answer:

```
"variables": {
  "dataDisks": {
    "caching": [
      None
      ReadOnly
      ReadWrite
    ]
  }
}

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

**NEW QUESTION: 201**

What should you implement to meet the disaster recovery requirements for the PaaS solution?

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

Answer: [\(SHOW ANSWER\)](#)

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

The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active geo-replication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy. The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.

Reference:

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

**NEW QUESTION: 202**

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

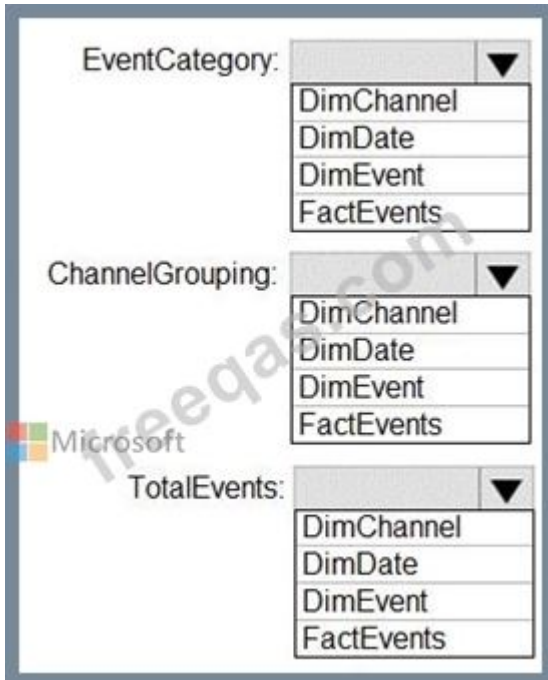
The data contains the following columns:

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

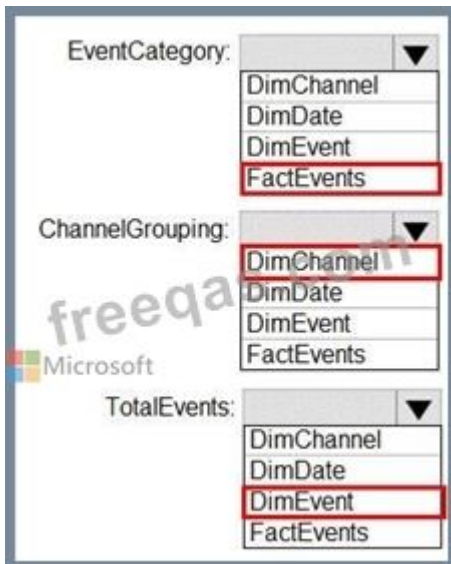
You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.

To which table should you add each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



**Answer:**



Reference:

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

### **NEW QUESTION: 203**

**HOTSPOT**

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

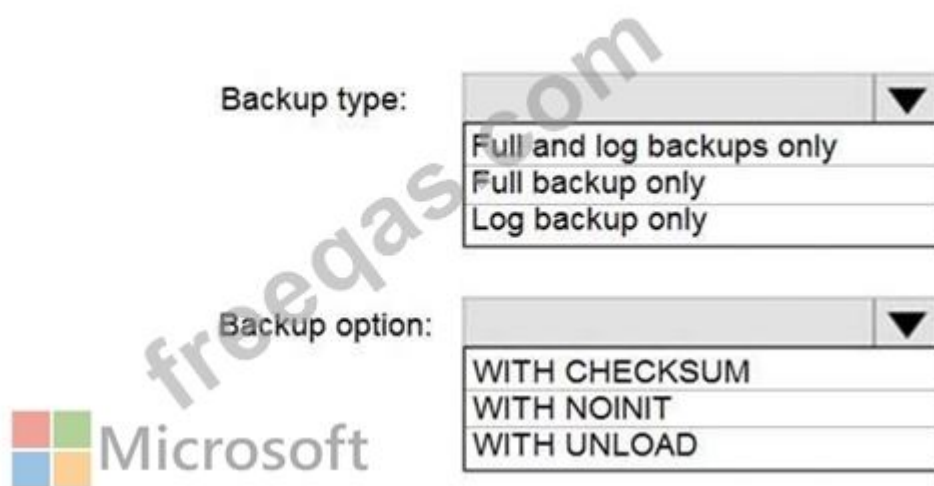
You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

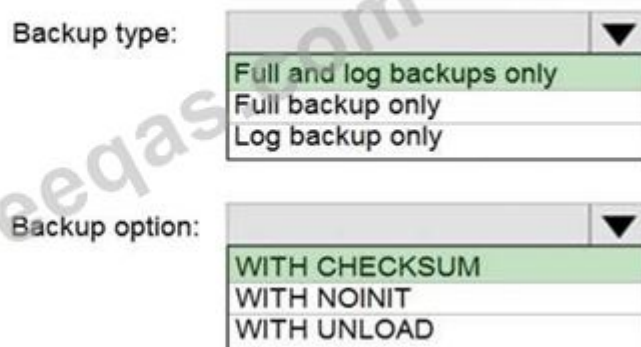
## Answer Area



The screenshot shows two dropdown menus. The first is labeled 'Backup type:' and has three options: 'Full and log backups only', 'Full backup only', and 'Log backup only'. The second is labeled 'Backup option:' and has three options: 'WITH CHECKSUM', 'WITH NOINIT', and 'WITH UNLOAD'. The Microsoft logo is visible on the left side of the image.

Answer:

Answer Area



The screenshot shows the same two dropdown menus as above. In this version, the correct answers are highlighted in green. The 'Backup type:' dropdown has 'Full and log backups only' selected. The 'Backup option:' dropdown has 'WITH CHECKSUM' selected.

Section: [none]

Explanation:

Box 1: Full and log backups only

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

Box 2: WITH CHECKSUM

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

Incorrect Answers:

NOINIT Indicates that the backup set is appended to the specified media set, preserving existing backup sets.

If a media password is defined for the media set, the password must be supplied. NOINIT is the default.

UNLOAD

Specifies that the tape is automatically rewound and unloaded when the backup is finished.

UNLOAD is the default when a session begins.

Reference:

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

**NEW QUESTION: 204**

You have SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

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

**Actions**

- Export a built-in policy.
- Create a custom policy based on a condition.
- Create a custom condition based on a built-in facet.
- View the policy history.
- Import a policy file.
- Run a policy evaluation.

**Answer Area**



**Answer:**

**Answer Area**



- Create a custom condition based on a built-in facet.
- Create a custom policy based on a condition.
- Run a policy evaluation.

- 1 - Create a custom condition based on a built-in facet.
- 2 - Create a custom policy based on a condition.
- 3 - Run a policy evaluation.

Reference:

<https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-based-management/>

**NEW QUESTION: 205**

You are evaluating the role assignments.

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

NOTE: Each correct selection is worth one point.

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

**Answer:**

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

Reference:

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

### NEW QUESTION: 206

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

You need to initiate a full backup of DB1 on SQL2.

Which statement should you run?

BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/

**A.** mycontainer/DB1.bak' with (Differential, STATS=5, COMPRESSION);

BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/

**B.** mycontainer/DB1.bak' with (COPY\_ONLY, STATS=5, COMPRESSION);

BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/

**C.** mycontainer/DB1.bak' with (File\_Snapshot, STATS=5, COMPRESSION);

BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/

**D.** mycontainer/DB1.bak' with (NoInit, STATS=5, COMPRESSION);

**Answer: (SHOW ANSWER)**

Section: [none]

Explanation:

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

Incorrect Answers:

A: Differential backups are not supported on secondary replicas. The software displays this error because the secondary replicas support copy-only database backups.

Reference:

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

### **NEW QUESTION: 207**

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

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

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

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

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

Does this meet the goal?

A. Yes

B. No

**Answer: A (LEAVE A REPLY)**

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

Reference:

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

### **NEW QUESTION: 208**

You have an Azure Data Lake Storage Gen2 container.

Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that. Users can read files in the container but cannot modify the files.

You need to design a data archiving solution that meets the following requirements:

New data is accessed frequently and must be available as quickly as possible.

Data that is older than five years is accessed infrequently but must be available within one second when requested.

Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.

Costs must be minimized while maintaining the required availability.

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

Five-year-old data:

	▼
Delete the blob.	
Move to archive storage.	
Move to cool storage.	
Move to hot storage.	

Seven-year-old data:

	▼
Delete the blob.	
Move to archive storage.	
Move to cool storage.	
Move to hot storage.	



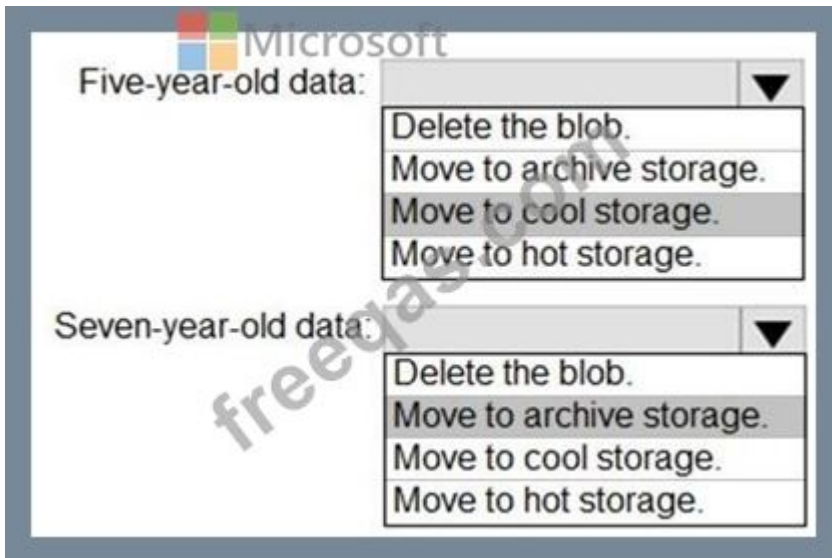
**Answer:**

Five-year-old data:		▼
Delete the blob.		
Move to archive storage.		
Move to cool storage.		
Move to hot storage.		
Seven-year-old data:		▼
Delete the blob.		
Move to archive storage.		
Move to cool storage.		
Move to hot storage.		



Explanation

Text, table Description automatically generated



#### Box 1: Move to cool storage

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

Short-term backup and disaster recovery

Older data not used frequently but expected to be available immediately when accessed Large data sets that need to be stored cost effectively, while more data is being gathered for future processing Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

#### Box 2: Move to archive storage

Example usage scenarios for the archive access tier include:

Long-term backup, secondary backup, and archival datasets

Original (raw) data that must be preserved, even after it has been processed into final usable form Compliance and archival data that needs to be stored for a long time and is hardly ever accessed Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

### **NEW QUESTION: 209**

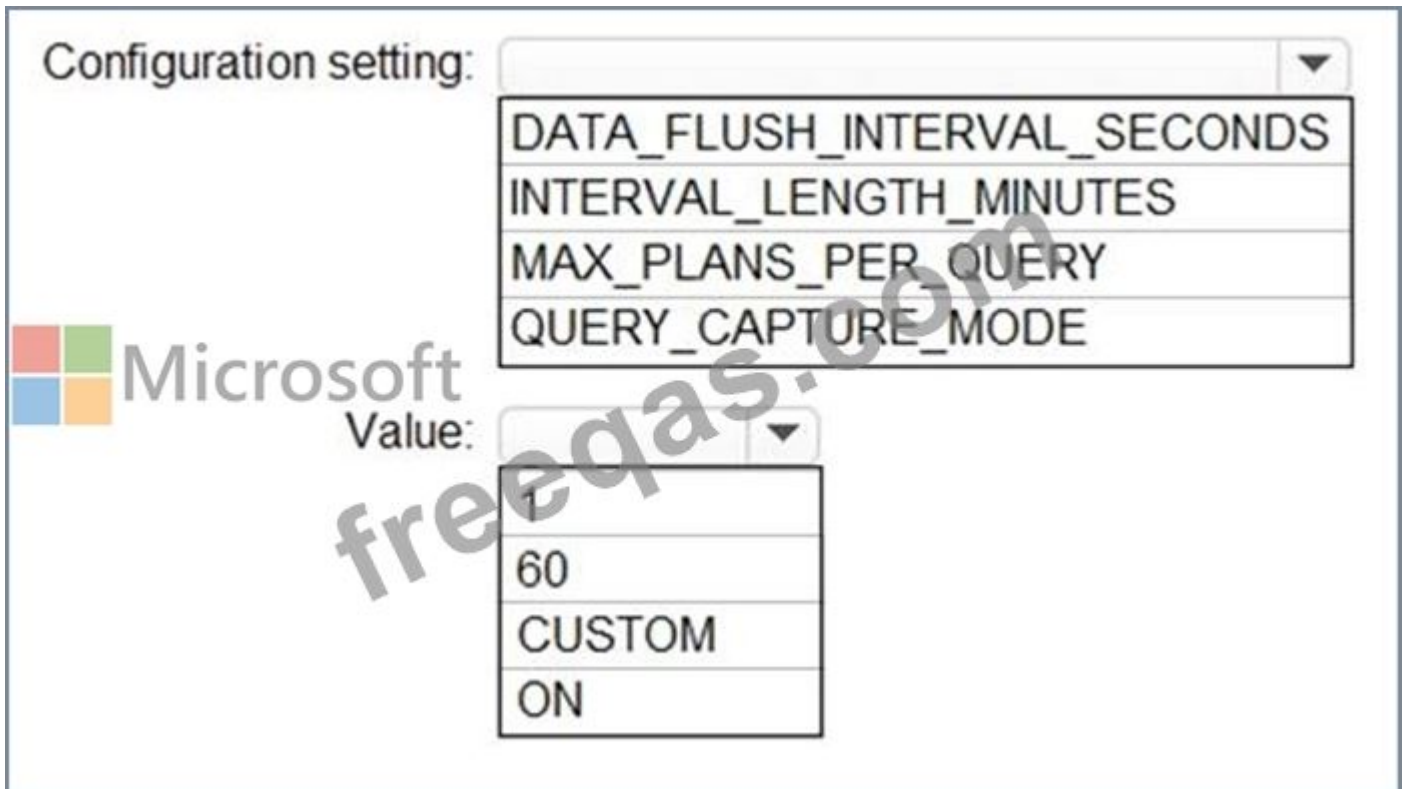
You have an Azure SQL database named db1 on a server named server1.

You use Query Performance Insight to monitor db1.

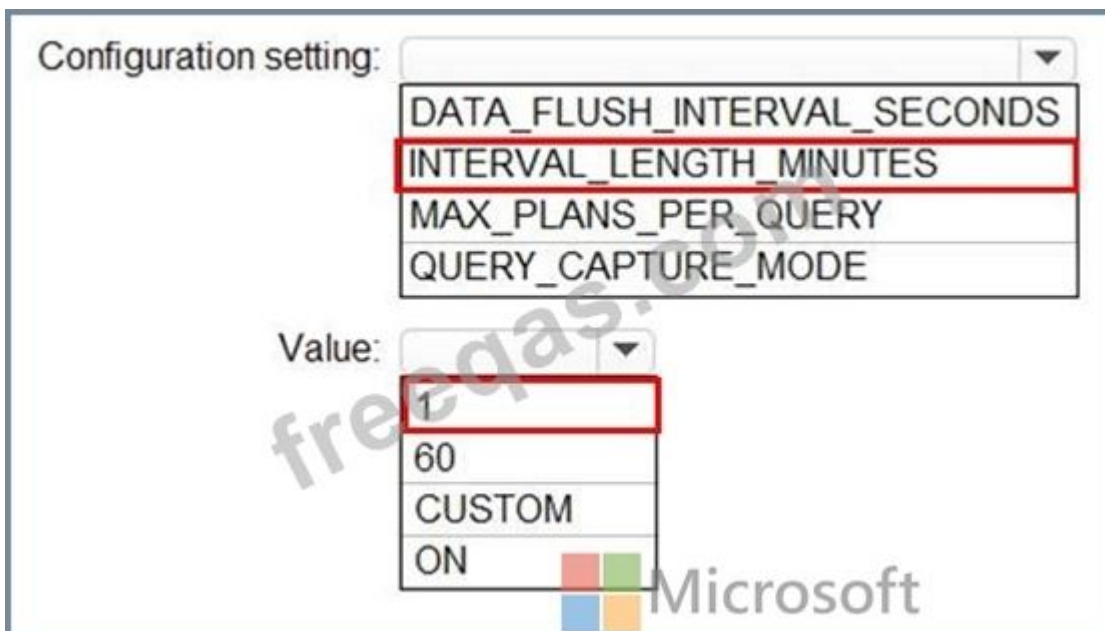
You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.

Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



**Answer:**



**NEW QUESTION: 210**

You have an Azure SQL logical server.

You run the following script.

```

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

```

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

**NOTE:** Each correct selection is worth one point

Statements	Yes	No
The <code>orders</code> table will allow only rows to be inserted.	<input type="radio"/>	<input type="radio"/>
To create additional tables in the Sales database, the <code>LEDGER = ON</code> parameter must be used.	<input type="radio"/>	<input type="radio"/>
To ensure that a timestamp is added to each row in the <code>orders</code> table, the <code>GENERATED ALWAYS</code>	<input type="radio"/>	<input type="radio"/>

**Answer:**

Statements	Yes	No
The <code>orders</code> table will allow only rows to be inserted.	<input checked="" type="radio"/>	<input type="radio"/>
To create additional tables in the Sales database, the <code>LEDGER = ON</code> parameter must be used.	<input type="radio"/>	<input checked="" type="radio"/>
To ensure that a timestamp is added to each row in the <code>orders</code> table, the <code>GENERATED ALWAYS</code>	<input type="radio"/>	<input checked="" type="radio"/>

### NEW QUESTION: 211

You have SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

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

## Actions

- Export a built-in policy.
- Create a custom policy based on a condition.
- Create a custom condition based on a built-in facet.
- View the policy history.
- Import a policy file.
- Run a policy evaluation.

## Answer Area

- Create a customer condition based on a built-in facet.
- Create a custom policy based on a condition.
- Run a policy evaluation.

## Answer:

- 1 - Create a customer condition based on a built-in facet.
- 2 - Create a custom policy based on a condition.
- 3 - Run a policy evaluation.

## Reference:

<https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-based-management/>

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!

PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

## NEW QUESTION: 212

### DRAG DROP

You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database.

A recent failure causes the master database to be lost.

You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.

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

Select and Place:

Actions	Answer Area
Add a certificate to an Azure key vault	
Enable Transparent Data Encryption (TDE)	
Encrypt a copy of the master key by using the service master key	
Turn on the TRUSTWORTHY property and the CLR property	
Attach the SSISDB database	
Open the master key for the SSISDB database	

Answer:

Actions	Answer Area
Add a certificate to an Azure key vault	Attach the SSISDB database
Enable Transparent Data Encryption (TDE)	Turn on the TRUSTWORTHY property and the CLR property
Encrypt a copy of the master key by using the service master key	Open the master key for the SSISDB database
Turn on the TRUSTWORTHY property and the CLR property	Encrypt a copy of the master key by using the service master key
Attach the SSISDB database	
Open the master key for the SSISDB database	

Explanation:

Step 1: Attach the SSISDB database

Step 2: Turn on the TRUSTWORTHY property and the CLR property

If you are restoring the SSISDB database to an SQL Server instance where the SSISDB catalog was never created, enable common language runtime (clr) Step 3: Open the master key for the SSISDB database Restore the master key by this method if you have the original password that was used to create SSISDB.

open master key decryption by password = 'LS1Setup!' --'Password used when creating SSISDB'  
Alter Master Key Add encryption by Service Master Key Step 4: Encrypt a copy of the mater key by using the service master key Reference:

<https://docs.microsoft.com/en-us/sql/integration-services/backup-restore-and-move-the-ssis-catalog> Implement a Secure Environment Testlet 1 This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Existing Environment

Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com.

All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

Database Environment

The sales department has the following database workload:

- \* An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

- \* A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

- \* An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability

group contains a database named ManufacturingSQLDb1 Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### Current Problems

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

#### Requirements

##### Planned Changes

Litware plans to implement the following changes:

- \* Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

- \* Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

- \* Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

- \* Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

- \* Migrate the SERVER1 databases to the Azure SQL Database platform.

##### Technical Requirements

Litware identifies the following technical requirements:

- \* Maintenance tasks must be automated.

- \* The 30 new databases must scale automatically.

- \* The use of an on-premises infrastructure must be minimized.

- \* Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

- \* All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

##### Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- \* Store encryption keys in Azure Key Vault.

- \* Retain backups of the PII data for two months.

- \* Encrypt the PII data at rest, in transit, and in use.

- \* Use the principle of least privilege whenever possible.

- \* Authenticate database users by using Active Directory credentials.

- \* Protect Azure SQL Database instances by using database-level firewall rules.

- \* Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

##### Business Requirements

Litware identifies the following business requirements:

- \* Meet an SLA of 99.99% availability for all Azure deployments.
- \* Minimize downtime during the migration of the SERVER1 databases.
- \* Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- \* Once all requirements are met, minimize costs whenever possible.

### NEW QUESTION: 213

You have the following Azure Resource Manager template.

```

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

```



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

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

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

A screenshot of a computer Description automatically generated with low confidence

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: 214

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

ProductID

ItemPrice

LineTotal  
Quantity  
StoreID  
Minute  
Month  
Hour  
Year  
Day

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

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

NOTE: Each correct selection is worth one point.

```
df.write
  .bucketBy (*)
  .partitionBy ("StoreID", "Hour")
  .range ("StoreID", "Year", "Month", "Day", "Hour")
  .sortBy ("Year", "Month", "Day", "Hour", "StoreID")
  .mode ("append")
  .parquet("/Purchases")
```

**Answer:**

```
df.write
  .bucketBy (*)
  .partitionBy ("StoreID", "Hour")
  .range ("StoreID", "Year", "Month", "Day", "Hour")
  .sortBy ("Year", "Month", "Day", "Hour", "StoreID")
  .mode ("append")
  .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: 215**

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails. You need to configure Azure Synapse Analytics to receive the data. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create an external data source for Azure Blob storage.	
Create a master key on database.	
Enable Transparent Data Encryption.	
Create the external table FactSalesOrderDetails.	
Load the data to a staging table.	
Create an external file format to map the parquet files.	

Answer:

Answer Area



Create a master key on database.

Create an external data source for Azure Blob storage.

Create an external file format to map the parquet files.

Create an external table FactSalesOrderDetails

- 1 - Create a master key on database.
- 2 - Create an external data source for Azure Blob storage.
- 3 - Create an external file format to map the parquet files.
- 4 - Create an external table FactSalesOrderDetails

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage

**NEW QUESTION: 216**

You have a SQL Server on Azure Virtual Machines instance named VM1 that hosts a database named DB1. You run the following query.

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

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

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

**Answer:**

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

**NEW QUESTION: 217**

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

You need to dynamically scale resources consumption based on workloads. The solution must minimize

downtime during scaling operations.

What should you use?

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

**Answer: A (LEAVE A REPLY)**

Section: [none]

Explanation:

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a

single server and share a set number of resources at a set price.

Reference:

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

### NEW QUESTION: 218

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

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

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

The query results include only [answer choice] in the csv/busfare folder.

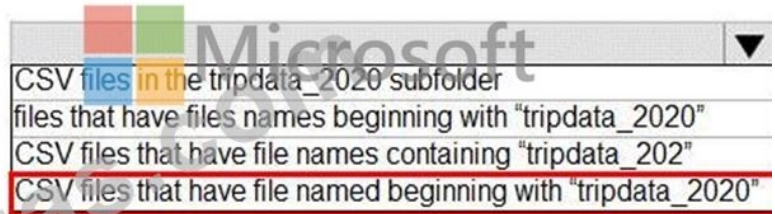
CSV files in the tripdata_2020 subfolder
files that have files names beginning with "tripdata_2020"
CSV files that have file names containing "tripdata_202"
CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is [answer choice] row.


a header
a data
an empty

Answer:

The query results include only [answer choice] in the csv/busfare folder.



The query assumes that the first row in a CSV file is [answer choice] row.



Reference:

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

### NEW QUESTION: 219

You have a burstable Azure virtual machine named VMI that hosts an instance of Microsoft SQL Server.

You need to attach an Azure ultra disk to VMI. The solution must minimize downtime on VMI. In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

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

**Answer Area**

Microsoft

**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** Microsoft

- 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: 220**

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named App1. App1 uses db1.

You need to enable Resource Governor for a App1. The solution must meet the following requirements:

App1 must be able to consume all available CPU resources.

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

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

**Actions**

- Create a plan.
- Create a classifier function in db1.
- Create a workload group.
- Create a classifier function in the master database.
- Create a resource pool that has the following configurations.
  - MAX\_CPU\_PERCENT = 100
  - MIN\_CPU\_PERCENT = 50

**Answer Area**

Microsoft

**Answer:**

## Answer Area

Create a resource pool that has the following configurations.....

Create a workload group.

Create a classifier function in the master database.



- 1 - Create a resource pool that has the following configurations.....
- 2 - Create a workload group.
- 3 - Create a classifier function in the master database.

### NEW QUESTION: 221

You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

- \* Clustered columnstore indexes
- \* Automatic tuning
- \* Change tracking
- \* PolyBase

You plan to migrate DB1 to an Azure SQL database.

What feature should be removed or replaced before DB1 can be migrated?

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

**Answer: B (LEAVE A REPLY)**

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

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

Incorrect Answers:

C: Change tracking is a lightweight solution that provides an efficient change tracking mechanism for applications. It applies to both Azure SQL Database and SQL Server.

D: Azure SQL Database and Azure SQL Managed Instance automatic tuning provides peak performance and stable workloads through continuous performance tuning based on AI and machine learning.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary> Implement a Secure Environment Question Set 1

### NEW QUESTION: 222

You have an Azure SQL managed instance.

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

What should you do?

- A. Generate an estimated execution plan.
- B. Generate an actual execution plan.
- C. Run sys.dm\_exec\_query\_plan\_scacs.
- D. Generate Live Query Statistics.

**Answer: C (LEAVE A REPLY)**

Reference:

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

**NEW QUESTION: 223**

You have the following Azure Resource Manager template.

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

```

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

NOTE: Each correct selection is worth one point.

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

**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"/>

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: 224

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

ProductID

ItemPrice

LineTotal

Quantity

StoreID

Minute

Month

Hour

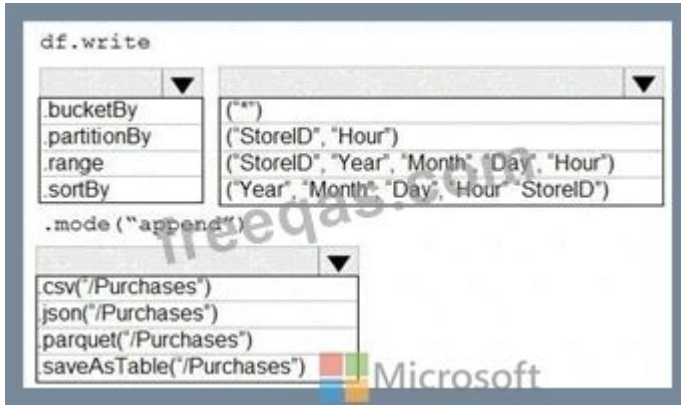
Year

Day

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

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

NOTE: Each correct selection is worth one point.



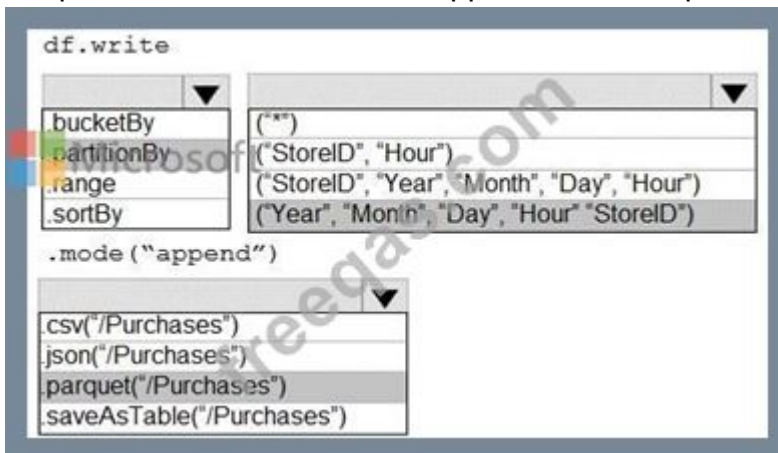
**Answer:**

df.write



Explanation

Graphical user interface, text, application Description automatically generated



Box 1: .partitionBy

Example:

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

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

```
parquet("/data/hive/warehouse/db_name.db/" + tableName)
```

Box 2: ("Year", "Month", "Day", "Hour", "StoreID")

Box 3: .parquet("/Purchases")

Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partitio>

**NEW QUESTION: 225**

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.

You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.

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

**Actions**

- Deploy Grafana to an Azure virtual machine.
- Build a `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file.
- Create Dropwizard counters in the application code.
- Create a data source in Azure Monitor.
- Configure the Databricks cluster to use the Databricks monitoring library.

**Answer Area**

Microsoft

**Answer:**

**Answer Area** Microsoft

- Configure your Azure Databricks cluster to use the Databricks monitoring library.
- Build the `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file
- Create Dropwizard counters in your application code.

- 1 - Configure your Azure Databricks cluster to use the Databricks monitoring library.
- 2 - Build the `spark-listeners-loganalytics-1.0-SNAPSHOT.jar` JAR file
- 3 - Create Dropwizard counters in your application code.

**NEW QUESTION: 226**

You have an Azure Synapse Analytics dedicated SQL pool.

You run `PDW_SHOWSPACEUSED('dbo.FactInternetSales');` and get the results shown in the

following table.

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

Which statement accurately describes the dbo.FactInternetSales table?

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

**Answer: (SHOW ANSWER)**

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>

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

**NEW QUESTION: 227**

You have an Azure subscription.

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

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

NOTE: Each correct selection is worth one point.

**Answer Area**

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

**Answer:**

**Answer Area**

```

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

```

Reference:

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

**NEW QUESTION: 228**

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

The data contains the following columns:

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

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.

To which table should you add each column? To answer, select the appropriate options in the

answer area.

NOTE:Each correct selection is worth one point.

EventCategory:   
DimChannel  
DimDate  
DimEvent  
FactEvents

ChannelGrouping:   
DimChannel  
DimDate  
DimEvent  
FactEvents

TotalEvents:   
DimChannel  
DimDate  
DimEvent  
FactEvents

**Answer:**

EventCategory:   
DimChannel  
DimDate  
DimEvent  
FactEvents

ChannelGrouping:   
DimChannel  
DimDate  
DimEvent  
FactEvents

TotalEvents:   
DimChannel  
DimDate  
DimEvent  
FactEvents

Explanation

Graphical user interface, application, table Description automatically generated



Box 1: FactEvents

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

Box 2: DimChannel

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

Box 3: DimEvent

Reference:

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

### NEW QUESTION: 229

You have an Azure SQL Database instance named DatabaseA on a server named Server1. You plan to add a new user named App1 to DatabaseA and grant App1 db\_datacenter permissions. App1 will use SQL Server Authentication.

You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.

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

Actions	Answer Area
On the master database, run CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;	
On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';	
On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];	
On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';	
On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];	

**Answer:**

Actions	Answer Area
On the master database, run CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;	On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';
On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';	On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];
On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];	On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];
On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';	
On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];	

**Explanation:**

Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!' Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:

```
CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
```

You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command.

Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database.

Here is some sample Transact-SQL that creates a user:

```
CREATE USER readonlyuser FROM LOGIN readonlylogin;
```

Step 3: On DatabaseA run ALTER ROLE db\_datareader ADD Member [App1]

Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to

the database via the db\_datareader role.

EXEC sp\_addrolemember 'db\_datareader', 'readonlyuser';

Reference:

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

### NEW QUESTION: 230

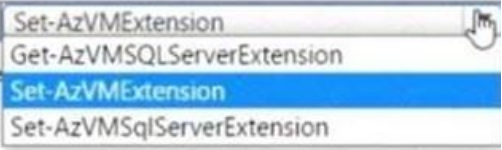

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL. You need to use PowerShell to enable and configure automated patching for SQL. The solution must include both SQL Server and Windows security updates.

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

NOTE: Each correct selection is worth one point.

#### Answer Area

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"  
  
Get-AzVM -ResourceGroupName "RG1" -Name "SQL" |  
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM  
-SQLManagementType
```



Microsoft

Answer:



Microsoft

Explanation

## Answer Area

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"  
  
Get-AzVM -ResourceGroupName "RG1" -Name "SQL1" | Set-AzVMExtension  
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM  
-SQLManagementType -Lightweight
```

### NEW QUESTION: 231

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

What should you include in the recommendation?

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

**Answer:** ([SHOW ANSWER](#))

Topic 1, Litware

Existing Environment

Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named `litwareinc.com`. All Azure subscriptions are associated to the `litwareinc.com` Azure AD tenant.

Database Environment

The sales department has the following database workload:

An on-premises named `SERVER1` hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

A logical server named `SalesSrv01A` contains a geo-replicated Azure SQL database named `SalesSQLDb1`. `SalesSQLDb1` is in an elastic pool named `SalesSQLDb1Pool`. `SalesSQLDb1` uses database firewall rules and contained database users.

An application named `SalesSQLDb1App1` uses `SalesSQLDb1`.

The manufacturing office contains two on-premises SQL Server 2016 servers named `SERVER2` and `SERVER3`. The servers are nodes in the same Always On availability group. The availability group contains a database named `ManufacturingSQLDb1`. Database administrators have two

Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### Current Problems

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

#### Requirements

##### Planned Changes

Litware plans to implement the following changes:

Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.

Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Migrate the SERVER1 databases to the Azure SQL Database platform.

##### Technical Requirements

Litware identifies the following technical requirements:

Maintenance tasks must be automated.

The 30 new databases must scale automatically.

The use of an on-premises infrastructure must be minimized.

Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

##### Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

Store encryption keys in Azure Key Vault.

Retain backups of the PII data for two months.

Encrypt the PII data at rest, in transit, and in use.

Use the principle of least privilege whenever possible.

Authenticate database users by using Active Directory credentials.

Protect Azure SQL Database instances by using database-level firewall rules.

Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

##### Business Requirements

Litware identifies the following business requirements:

Meet an SLA of 99.99% availability for all Azure deployments.

Minimize downtime during the migration of the SERVER1 databases.

Use the Azure Hybrid Use Benefits when migrating workloads to Azure.  
Once all requirements are met, minimize costs whenever possible.

**NEW QUESTION: 232**

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1.

You need to encrypt the data in Col1.

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

**Actions**

**Answer Area**

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

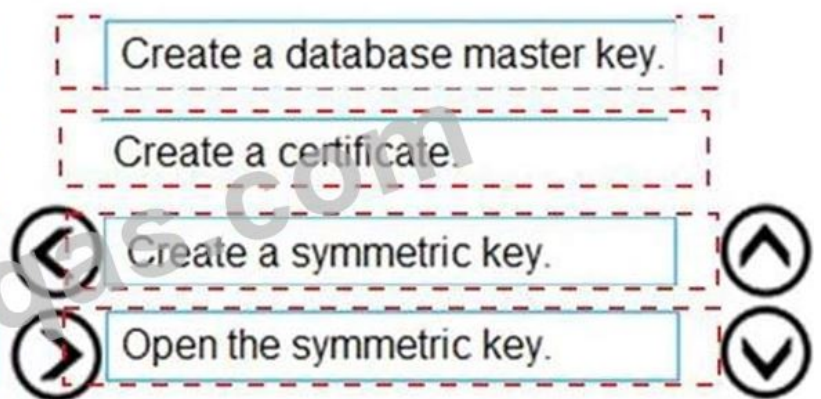


Answer:

**Actions**

**Answer Area**

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



Explanation

Create a database master key.

Create a certificate.

Create a symmetric key.

Open the symmetric key.



Table Description automatically generated

Reference:

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

### NEW QUESTION: 233

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

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

Which resource type should you use to create the alert?

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

**Answer: (SHOW ANSWER)**

Section: [none]

Explanation:

There are resource types related to application code, compute infrastructure, networking, storage + databases.

You can deploy up to 800 instances of a resource type in each resource group.

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

Reference:

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

### NEW QUESTION: 234

You have an Azure SQL managed instance.

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

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

RESTORE **Microsoft** DB1 FROM

DATABASE  
FILE  
LOG

DISK = N''\NAS01\SSQLBackups\DB1.bak';  
TAPE = N''\Tape0''  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak''

**Answer:**

RESTORE **Microsoft** DB1 FROM

DATABASE  
FILE  
LOG

DISK = N''\NAS01\SSQLBackups\DB1.bak';  
TAPE = N''\Tape0''  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak''

Explanation

Text Description automatically generated

RESTORE **Microsoft** DB1 FROM

DATABASE  
FILE  
LOG

DISK = N''\NAS01\SSQLBackups\DB1.bak';  
TAPE = N''\Tape0''  
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak''

### NEW QUESTION: 235

You have an Azure SQL database named DB1.

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

During a database audit, you discover that none of the fields contain more than 150 characters.

You need to ensure that you can apply page compression to Table1.

What should you do?

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

**Answer: D (LEAVE A REPLY)**

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: 236

You have the following Transact-SQL query.

```

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

```

FROM sys.database\_files:

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

- A. ColumnA
- B. ColumnB
- C. ColumnC
- D. ColumnD

**Answer:** ([SHOW ANSWER](#))

Example:

Free space for the file in the below query result set will be returned by the FreeSpaceMB column.

```

SELECT DB_NAME() AS DbName,
name AS FileName,
type_desc,
size/128.0 AS CurrentSizeMB,
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB

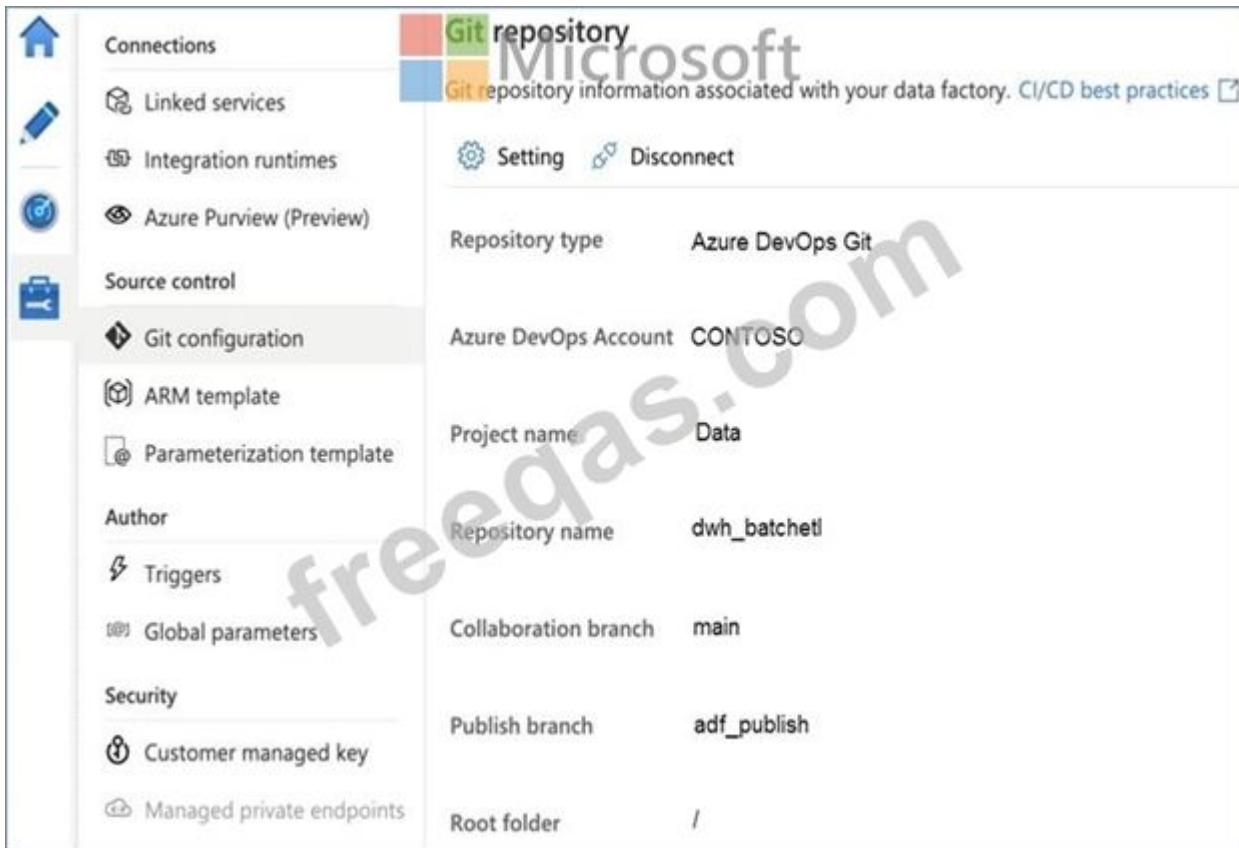
```

FROM sys.database\_files WHERE type IN (0,1); Reference:

<https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/>

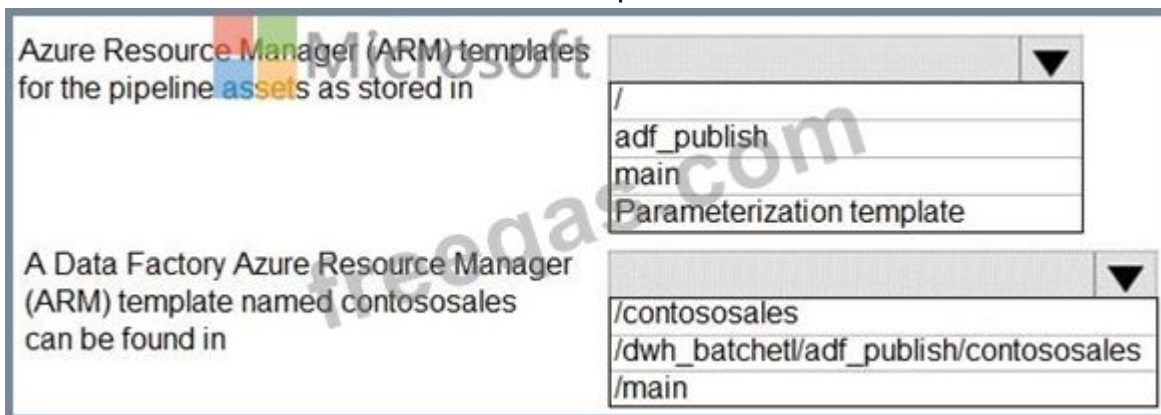
### NEW QUESTION: 237

You configure version control for an Azure Data Factory instance as shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

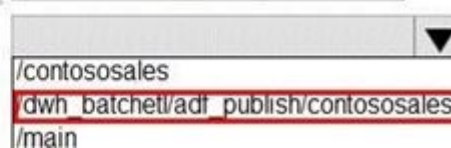


**Answer:**

Azure Resource Manager (ARM) templates for the pipeline assets as stored in



A Data Factory Azure Resource Manager (ARM) template named contososales can be found in



Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/source-control>

**NEW QUESTION: 238**

You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

Clustered columnstore indexes

Automatic tuning

Change tracking

PolyBase

You plan to migrate DB1 to an Azure SQL database.

What feature should be removed or replaced before DB1 can be migrated?

**A.** Clustered columnstore indexes

**B.** PolyBase

**C.** Change tracking

**D.** Automatic tuning

**Answer: B (LEAVE A REPLY)**

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

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

Reference:

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

**NEW QUESTION: 239**

Which audit log destination should you use to meet the monitoring requirements?

**A.** Azure Storage

**B.** Azure Event Hubs

## C. Azure Log Analytics

**Answer: C (LEAVE A REPLY)**

### 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>

### Topic 2, Litware

#### Existing Environment

##### Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

##### Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

##### Database Environment

The sales department has the following database workload:

An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1 Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

##### Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

##### Current Problems

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

##### Requirements

## Planned Changes

Litware plans to implement the following changes:

Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Migrate the SERVER1 databases to the Azure SQL Database platform.

## Technical Requirements

Litware identifies the following technical requirements:

Maintenance tasks must be automated.

The 30 new databases must scale automatically.

The use of an on-premises infrastructure must be minimized.

Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

## Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

Store encryption keys in Azure Key Vault.

Retain backups of the PII data for two months.

Encrypt the PII data at rest, in transit, and in use.

Use the principle of least privilege whenever possible.

Authenticate database users by using Active Directory credentials.

Protect Azure SQL Database instances by using database-level firewall rules.

Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

## Business Requirements

Litware identifies the following business requirements:

Meet an SLA of 99.99% availability for all Azure deployments.

Minimize downtime during the migration of the SERVER1 databases.

Use the Azure Hybrid Use Benefits when migrating workloads to Azure.

Once all requirements are met, minimize costs whenever possible.

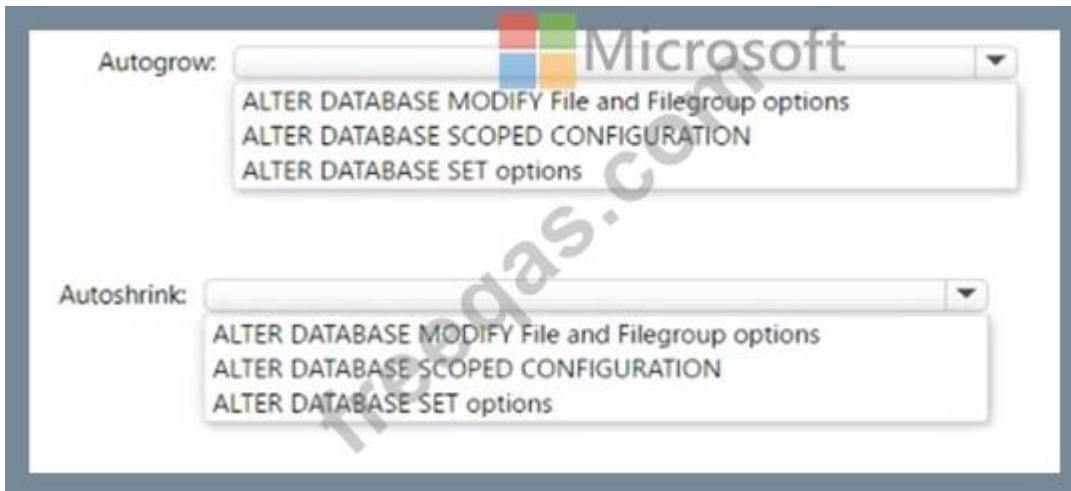
## **NEW QUESTION: 240**

You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1.

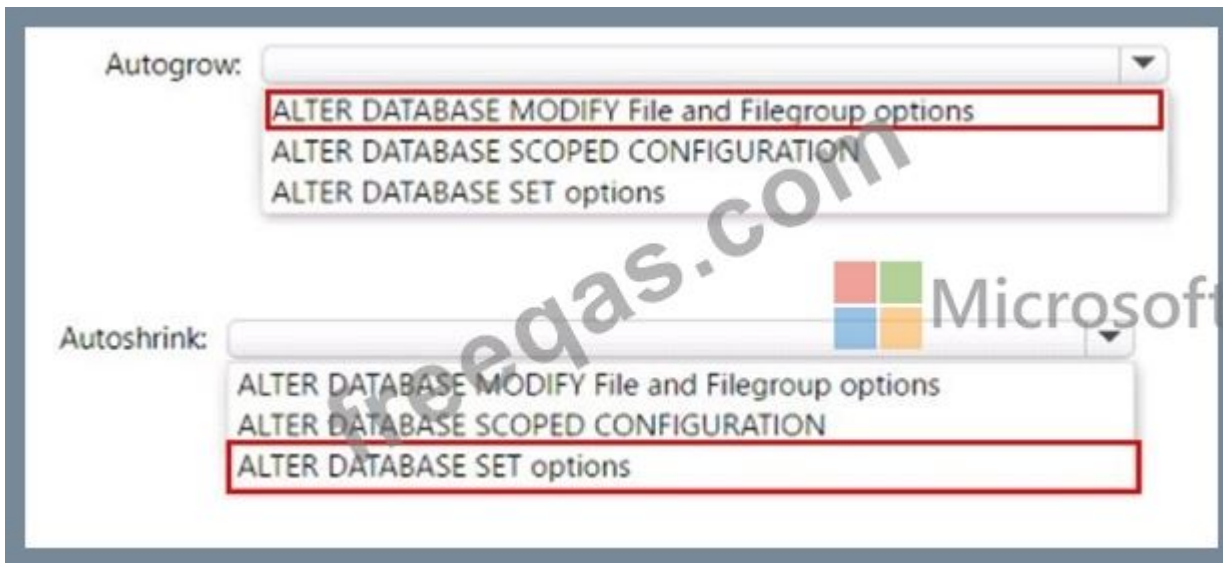
You need to configure the autogrow and autoshrink settings for DB1.

Which statements should you use? To answer, select the appropriate options in the answer area a.

NOTE: Each correct selection is worth one point.



**Answer:**



**NEW QUESTION: 241**

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

### Configure policies ✕

SQL server

#### Point in Time Restore Configuration

Configure PiTR backup retention ▼ Days

---

#### Long-term Retention Configurations

Weekly LTR Backups ⓘ

How long would you like weekly backups to be kept?

6 Week(s) ▼

---

Monthly LTR Backups ⓘ

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

12 Month(s) ▼

---

Yearly LTR Backups ⓘ

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

Week 2 ▼

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

10 Year(s) ▼

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

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

February 29, 2020

March 7, 2020

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

NOTE: Each correct selection is worth one point.

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

- 6 weeks
- 12 months
- 10 years

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

- 6 weeks
- 12 months
- 10 years

**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

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!  
PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com DP-300 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

**NEW QUESTION: 242**

You have the following Azure Resource Manager template.

```

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

```

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

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

Answer:



## Statements

Yes

No

The template deploys a serverless Azure SQL database.

The template deploys a database to an Azure SQL Database managed instance.

The pricing tier of the database deployment is based on DTUs.

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: 243

You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1.

You need to identify and repair any physical or logical corruption in DB1. The solution must meet the following requirements:

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

How should you complete the command? To answer, select the appropriate options in the answer area NOTE:

Each correct selection is worth one point.

Answer Area

DBCC CHECK [DB1],

REPAIR\_REBUILD

NOINDEX

REPAIR\_ALLOW\_DATA\_LOSS

REPAIR\_FAST

REPAIR\_REBUILD

WITH

PHYSICAL\_ONLY;

EXTENDED\_LOGICAL\_CHECKS;

PHYSICAL\_ONLY;

TABLOCK;

Answer:

Answer Area

DBCC CHECK [DB1],

REPAIR\_REBUILD

NOINDEX

REPAIR\_ALLOW\_DATA\_LOSS

REPAIR\_FAST

REPAIR\_REBUILD

) WITH

PHYSICAL\_ONLY;

EXTENDED\_LOGICAL\_CHECKS;

PHYSICAL\_ONLY;

TABLOCK;

Explanation

Answer Area

DBCC CHECK [DB1],

REPAIR\_REBUILD

) WITH

PHYSICAL\_ONLY;



**NEW QUESTION: 244**

You have an Azure Synapse Analytics Apache Spark pool named Pool1.

You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.

You need to load the files into the tables. The solution must maintain the source data types.

What should you do?

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

**Answer: B (LEAVE A REPLY)**

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools. Serverless SQL pool enables you to query data in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.

To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.

The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files>

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

**Valid DP-300 Dumps** shared by PrepPdf.com for Helping Passing DP-300 Exam!

PrepPdf.com now offer the **newest DP-300 exam dumps**, the PrepPdf.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest**

PrepPdf.com DP-300 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/DP-300-prepaway-exam-dumps.html> (343 Q&As Dumps,

**40%OFF Special Discount: Exam-Tests**)