

## Microsoft.DP-300.v2026-01-29.q198

<b>Exam Code:</b>	DP-300
<b>Exam Name:</b>	Administering Relational Databases on Microsoft Azure
<b>Certification Provider:</b>	Microsoft
<b>Free Question Number:</b>	198
<b>Version:</b>	v2026-01-29
<b># of views:</b>	108
<b># of Questions views:</b>	2120
<a href="https://www.freecram.net/torrent/Microsoft.DP-300.v2026-01-29.q198.html">https://www.freecram.net/torrent/Microsoft.DP-300.v2026-01-29.q198.html</a>	

### NEW QUESTION: 1

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

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

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

NOTE: Each correct selection is worth one point.

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

**Answer: (SHOW ANSWER)**

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

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

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

Incorrect Answers:

A: High - The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.

Reference:

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

### NEW QUESTION: 2

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

To improve the concurrency of tempdb, try the following methods:

- \* Increase the number of data files in tempdb to maximize disk bandwidth and reduce contention in allocation structures.

- \* Etc.

Note: Symptoms On a server that is running Microsoft SQL Server, you notice severe blocking when the server is experiencing a heavy load. Dynamic Management Views

[`sys.dm_exec_request` or `sys.dm_os_waiting_tasks`] indicates that these requests or tasks are waiting for tempdb resources. Additionally, the wait type is `PAGELATCH_UP`, and the wait resource points to pages in Tempdb.

Reference:

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

### **NEW QUESTION: 3**

Hotspot Question

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.

## Answer Area

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:

Answer Area Microsoft

ALTER DATABASE [Db1]

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

GO

ALTER DATABASE [Db1]

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

GO

Explanation:

Box 1: SET QUERY\_STORE = ON (OPERATION MODE = READ\_WRITE);

As a pre-requisite, Query Store must be turned on.

Box 2: SET AUTOMATIC\_TUNING (FORCE\_LAST\_GOOD\_PLAN = ON)

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

ALTER DATABASE current SET AUTOMATIC\_TUNING (FORCE\_LAST\_GOOD\_PLAN = ON)

Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option.

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

Reference:

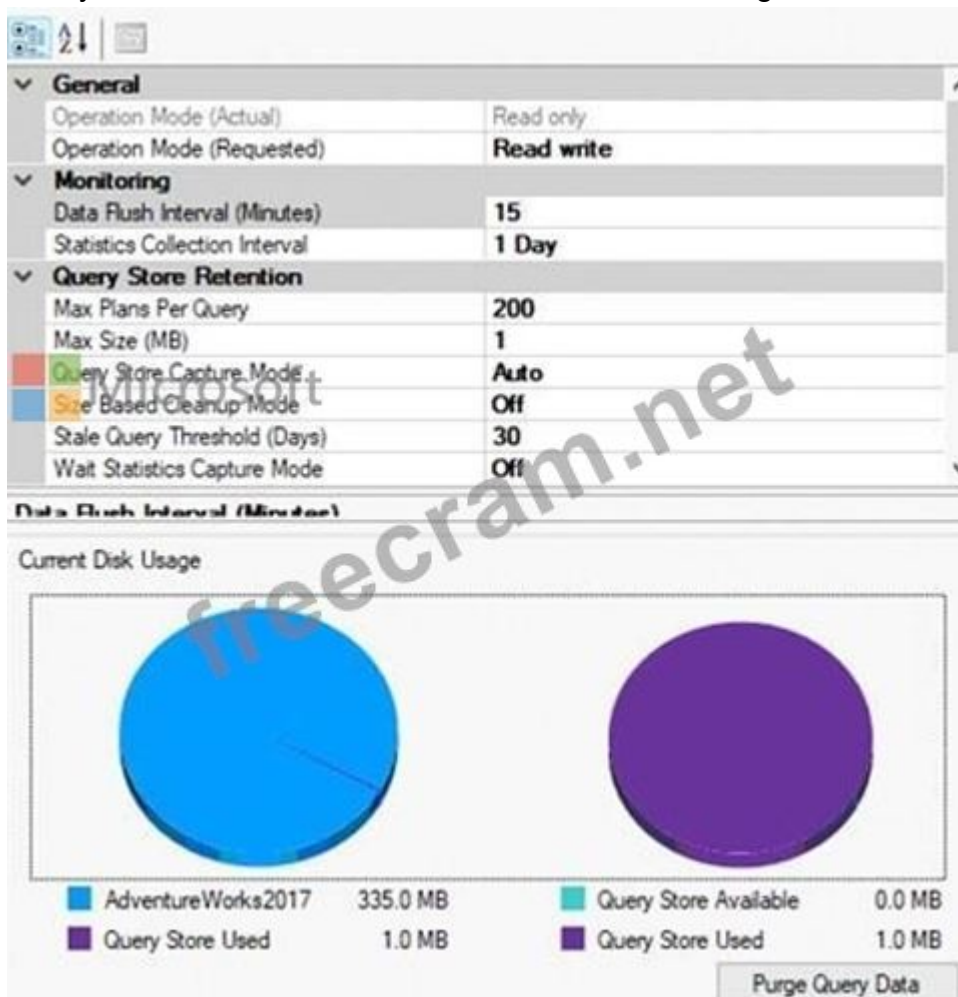
<https://docs.microsoft.com/en-us/sql/relational-databases/automatic-tuning/automatic-tuning>

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

## NEW QUESTION: 4

Hotspot Question

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.



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.

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

all
none of the
a selective set of

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)

**Answer:**

The screenshot shows the same question as above, but with the correct answers highlighted in green. In the first dropdown menu, "none of the" is selected. In the second dropdown menu, "Max Size (MB)" is selected.

**Explanation:**

The first answer is "none of the" because the Operation Mode (actual) is "Read only" so it is not retaining any query for evaluation due to the Query Store Available is 0.0 MB.

When the actual state is read-only, use the readonly\_reason column to determine the root cause.

Typically, you find that Query Store transitioned to read-only mode because the size quota was exceeded. In that case, the readonly\_reason is set to 65536 Consider the following steps to switch Query Store to read-write mode and activate data collection:

\* Increase the maximum storage size by using the MAX\_STORAGE\_SIZE\_MB option of ALTER DATABASE or

\* Clean up Query Store data

<https://learn.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store?view=sql-server-ver16#Verify>

**NEW QUESTION: 5**

**SIMULATION**

You need to configure db1 to pause automatically after one hour of inactivity.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

Azure SQL elastic pool pausing feature

An elastic pool vCore provisioned server can be paused. You can do this by changing the compute tier of the elastic pool to Serverless. In the serverless compute tier, the database is automatically paused when it is inactive for a period of time that you specify. The database is automatically resumed when the next login or other activity occurs.

Part 1: Change to the serverless compute tier

To change the compute tier of an elastic pool, you can follow these steps:

Step 1: Go to the Azure portal and sign in to your account.

Step 2: Select the Azure SQL resource that contains the elastic pool.

Step 3: In the left navigation menu, select Elastic Pools.

Step 4: Select the elastic pool that you want to change the compute tier for.

Step 5: In the Settings section, select Compute Tier.

Step 6: Select Serverless.

Step 7: Click Save.

Part 2: Configure the Auto-pause setting.

Step 1: Select db1.

Once you have changed the compute tier to serverless, the database will be paused after the idle time that you specified. You can view the idle time in the Auto-pause delay setting.

Step 2: Open database settings.

Step 3: Change the Auto-pause delay setting to 1 hour.

Reference:

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

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

**NEW QUESTION: 6**

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

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

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Column1

B. Column2

- C. Column3
- D. Column4
- E. Column5

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

Always Encrypted isn't supported for the columns using one of the following data types: xml, timestamp, rowversion, image, ntext, text, sql\_variant, hierarchyid, geography, geometry, alias, user-defined types.

### **NEW QUESTION: 7**

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

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: 8**

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](#))**

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

### **NEW QUESTION: 9**

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

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

What should you do?

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

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

### NEW QUESTION: 10

Hotspot Question

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	<input type="button" value="▼"/>	DB1	FROM	<input type="button" value="▼"/>
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	<input type="button" value="▼"/>	DB1	FROM	<input type="button" value="▼"/>
DATABASE				DISK = N'\\NAS01\SQLBackups\DB1.bak';
FILE				TAPE = N'\\Tape0'
LOG				URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Explanation:

Syntax for Azure SQL Managed Instance for the restore command:

--To Restore an Entire Database from a Full database backup (a Complete Restore):

RESTORE DATABASE { database\_name | @database\_name\_var }

FROM URL = { 'physical\_device\_name' | @physical\_device\_name\_var } [ ,...n ]

[;]

Arguments

DATABASE

Specifies the target database.

FROM URL

Specifies one or more backup devices placed on URLs that will be used for the restore operation.

The URL format is used for restoring backups from the Microsoft Azure storage service.

Reference:

[https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql?](https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql?view=azuresqldb-mi-current&preserve-view=true)

[view=azuresqldb-mi-current&preserve-view=true](https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql?view=azuresqldb-mi-current&preserve-view=true)

### NEW QUESTION: 11

You have an Azure SQL database named SQLDb1 that contains the resources shown in the following table.

Name	Type
Table1	Rowstore table
Column1	Nvarchar(max)

Column1 contains JSON data.

You need to compress Column1. The solution must minimize the amount of storage used.

What should you use?

- A. the COMPRESS() function
- B. columnstore archive compression
- C. row compression
- D. columnstore compression

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

<https://learn.microsoft.com/es-es/sql/t-sql/functions/compress-transact-sql?view=sql-server-ver16>

### NEW QUESTION: 12

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: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.

Does this meet the goal?

- A. Yes
- B. No

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

This procedure is necessary and correct, as there is no option in Azure SQL Database to directly restore a database.

Reference:

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

### NEW QUESTION: 13

Hotspot Question

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 Area

EventCategory:  ▼

DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping:  ▼

DimChannel
DimDate
DimEvent
FactEvents

TotalEvents:  ▼

DimChannel
DimDate
DimEvent
FactEvents

Answer:

## Answer Area

EventCategory:  ▼  
Microsoft  
DimChannel  
DimDate  
DimEvent  
FactEvents

ChannelGrouping:  ▼  
DimChannel  
DimDate  
DimEvent  
FactEvents

TotalEvents:  ▼  
DimChannel  
DimDate  
DimEvent  
FactEvents

Explanation:

Box 1: DimEvent

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: FactEvents

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

Reference:

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

## NEW QUESTION: 14

### SIMULATION

You need to configure a monthly backup of db2. The monthly backups must be retained for five years, and then deleted.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

### Answer:

Change automated backup settings for Azure SQL Database

You can modify automated backup settings for Azure SQL Database, such as the short-term retention policy and the backup storage redundancy option that's used for backups.

To change differential backup frequency, select 12 Hours or 24 hours from the drop-down menu under Differential backup frequency.

Step 1: Go to the logical server in Azure with the databases whose retention period you want to change.

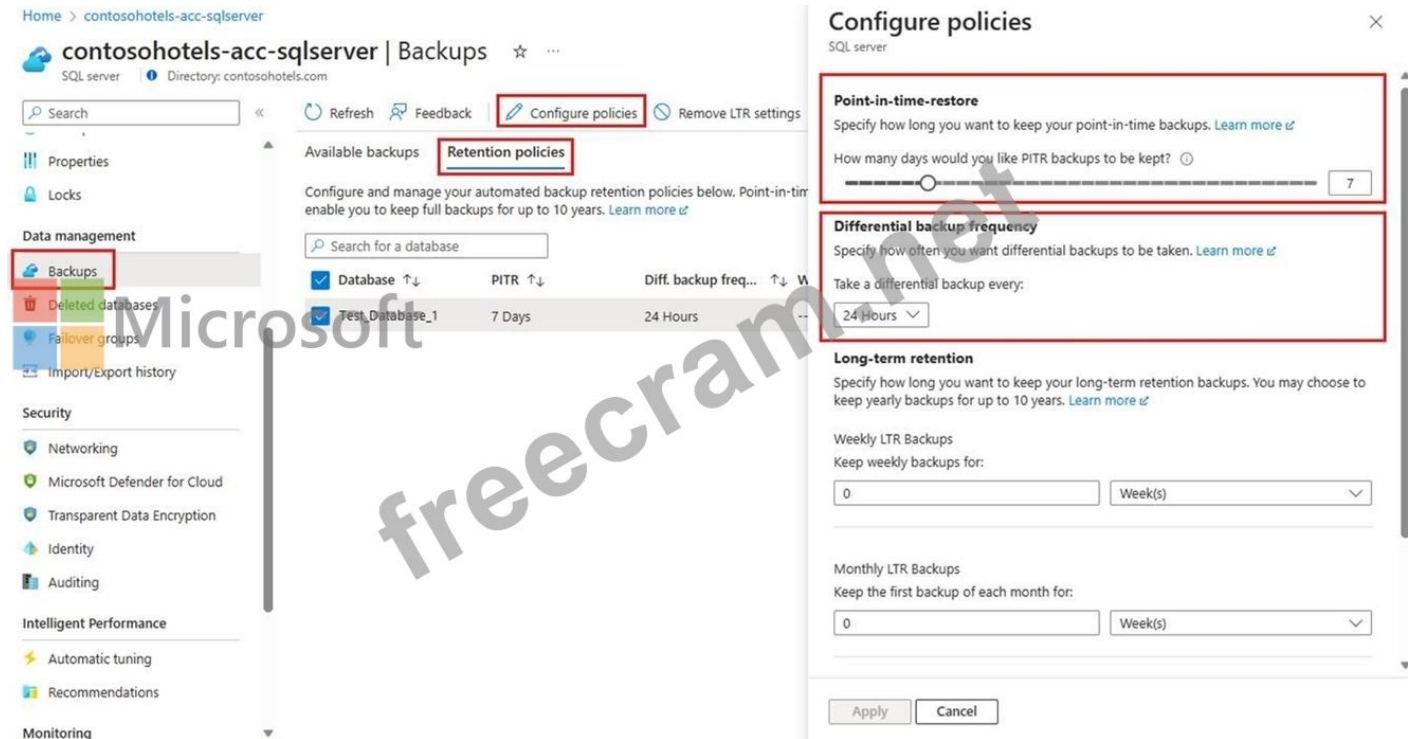
Step 2: Select Backups on the left pane, and then select the Retention policies tab.

Step 3: Select the databases for which you want to change the PITR backup retention. [Select db2]

Step 4: Select Configure policies from the action bar.

Step 5: To change the retention period for point-in-time restore backups, use the slider under Point-in-time restore. [Skip]

Step 6: To change the PITR (point-in-time recovery) backup retention period or the differential backup frequency for active databases by using the Azure portal: [Skip]



Step 7: In Long-term retention, Monthly LTR Backups, Set Keep the first backup of each month for: 5 Years Question: The monthly backups must be retained for five years, and then deleted.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/automated-backups-change-settings>

### NEW QUESTION: 15

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

Answer: B (LEAVE A REPLY)

### NEW QUESTION: 16

You have an Azure SQL database named DB1 that contains a nonclustered index named index1.

End users report slow queries when they use index1.

You need to identify the operations that are being performed on the index.

Which dynamic management view should you use?

- A. Sys.dm\_exec\_query\_plan\_stats
- B. Sys.dm\_db\_index\_physical\_stats
- C. Sys.dm\_db\_index\_operational\_stats
- D. Sys.dm\_db\_index\_usage\_stats

**Answer: (SHOW ANSWER)**

Every individual seek, scan, lookup, or update on the specified index by one query execution is counted as a use of that index and increments the corresponding counter in this view. Information is reported both for operations caused by user-submitted queries, and for operations caused by internally generated queries, such as scans for gathering statistics.

<https://learn.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-index-usage-stats-transact-sql?view=sql-server-ver16#remarks>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**)

**Special Discount Code: freecram)**

#### **NEW QUESTION: 17**

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. Azure Database Migration Service

**Answer: (SHOW ANSWER)**

When you migrate databases to Azure by using Azure Database Migration Service, you can do an offline or an online migration. With an offline migration, application downtime starts when the migration starts. With an online migration, downtime is limited to the time to cut over at the end of migration. We suggest that you test an offline migration to determine whether the downtime is acceptable; if not, do an online migration.

Reference:

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

#### **NEW QUESTION: 18**

Drag and Drop Question

You have a database named db1.

The log for db1 contains the following entry.

```
Date 10/5/2021 10:57:08 AM
Log SQL Server (Current - 10/5/2021 11:26:00 AM)
Sources spid1595
Message
The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'
```

You need to ensure that db1 can process transactions.

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**

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

**Answer Area**

- 
- 
- 

**Answer:**

**Actions**

- Shrink the transaction log file.
- Shrink db1.

**Answer Area**

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

**NEW QUESTION: 19**

You have an instance of SQL Server on Azure Virtual Machines named VM1.

You plan to schedule a SQL Server Agent job that will rebuild indexes of the databases hosted on VM1.

You need to configure the account that will be used by the agent. The solution must use the principle of least privilege.

Which operating system user right should you assign to the account?

- A. Increase scheduling priority
- B. Log on as a service
- C. Log on as a batch job
- D. Profile system performance

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

### NEW QUESTION: 20

Drag and Drop Question

You have an Azure subscription that contains a virtual network named VNet1 and two SQL Server on Azure Virtual Machines instances named SQLVM1 and SQLVM2. SQLVM1 and SQLVM2 are connected to VNet1.

You plan to deploy a domain-independent availability group that will be hosted on SQLVM1 and SQLVM2.

You create the required security principals.

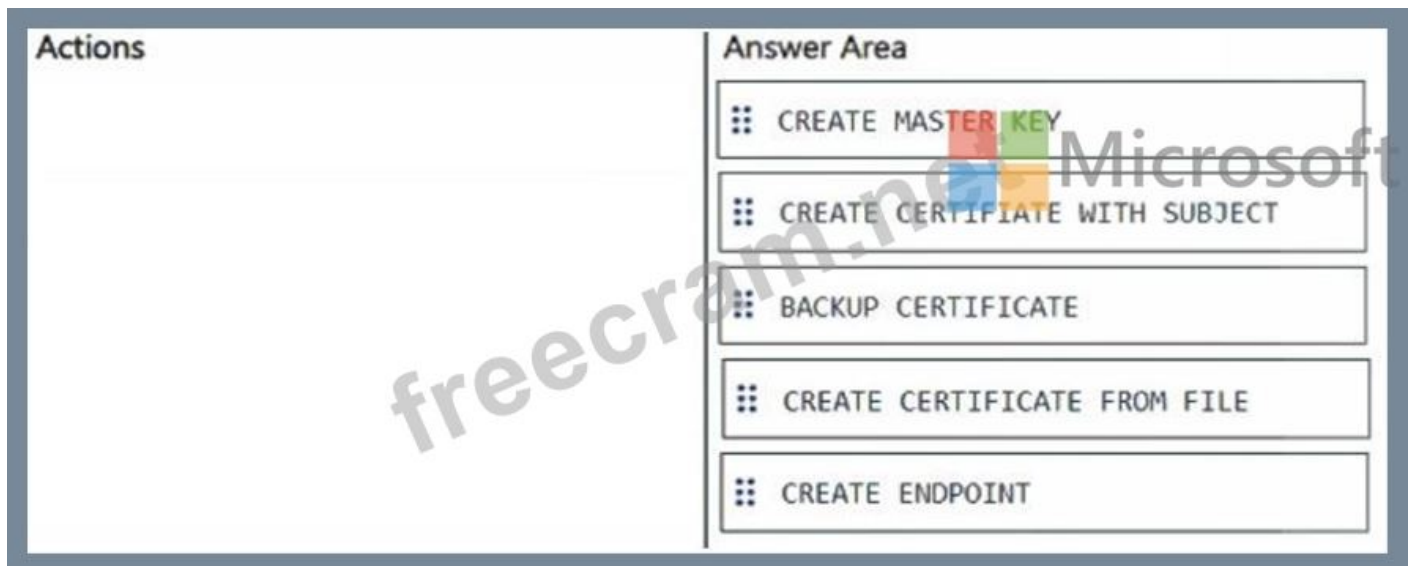
You need to complete the deployment.

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

Actions	Answer Area
CREATE CERTIFICATE FROM FILE	
CREATE ENDPOINT	
CREATE MASTER KEY	
BACKUP CERTIFICATE	
CREATE CERTIFICATE WITH SUBJECT	

*freecram.net*  
Microsoft

**Answer:**



## NEW QUESTION: 21

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event\_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements

Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a backup solution to restore DB3. The solution must meet the availability requirements.

Which type of backup should you use?

- A. differential
- B. transaction log
- C. long-term retention (LTR)
- D. point-in-time restore (PITR)

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

<https://azure.microsoft.com/en-us/blog/azure-sql-database-point-in-time-restore>

## NEW QUESTION: 22

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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.

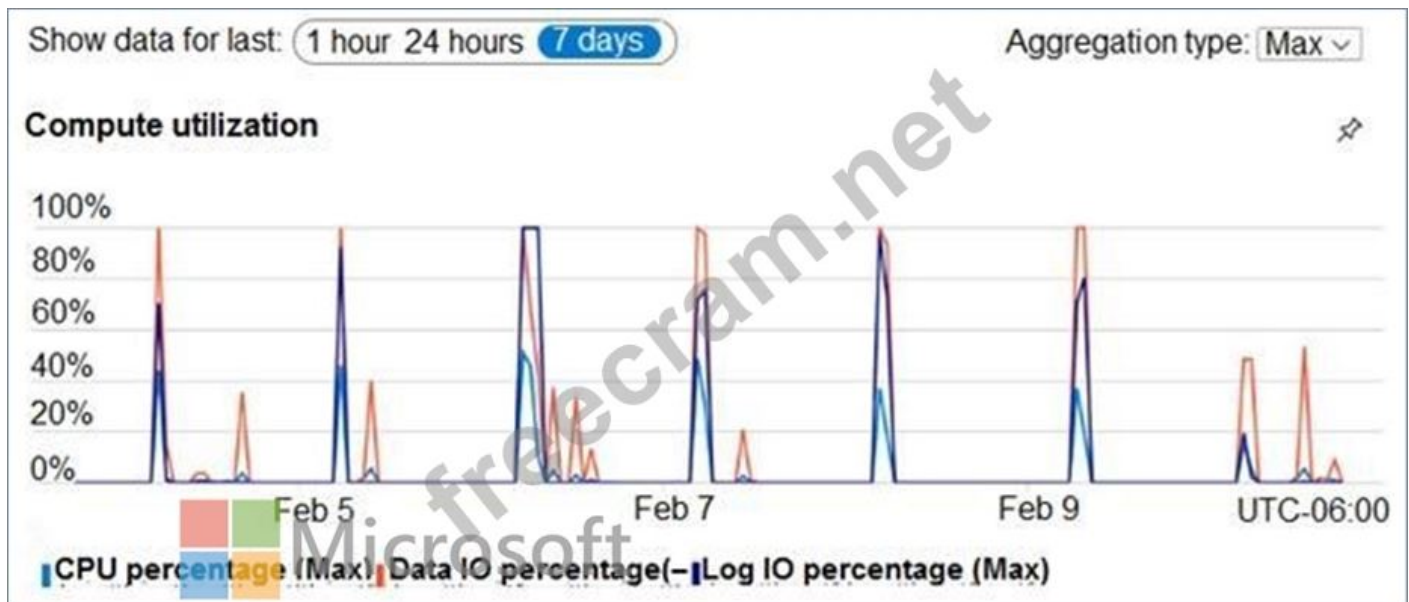
Use a single dashboard to review security and audit data for all the PaaS databases.

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.

Number of role assignments for this subscription: 15 / 2000

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

2 items (2 Groups)

Name	Type	Role	Scope
<input type="checkbox"/> DBAGroup1	Group	Contributor	This resource
<input type="checkbox"/> DBAGroup2	Group	SQL DB Contributor	This resource

You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements.

What should you do?

- Create an Azure Monitor alert rule that has a static threshold and assign the alert rule to an action group.
- Add a diagnostic setting that logs QueryStoreRuntimeStatistics and streams to an Azure event hub.
- Add a diagnostic setting that logs Timeouts and streams to an Azure event hub.
- Create an Azure Monitor alert rule that has a dynamic threshold and assign the alert rule to an action group.

**Answer: (SHOW ANSWER)**

Notify administrators when a PaaS database has a higher than average CPU usage. With Dynamic Thresholds, we use a unique ML technology to identify the patterns and come up with a single alert rule that has the right thresholds and accounts for seasonality patterns such as hourly, daily, or weekly.

Reference:

<https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/>

**NEW QUESTION: 23**

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1. App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event\_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

#### Requirements

#### Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

#### Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

#### Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

#### Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

#### Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to identify the event\_file target for monitoring DB3 after the migration to Azure SQL Database. The solution must meet the management requirements.

What should you use as the event\_file target?

- A. an Azure Blob Storage container
- B. an Azure Files share
- C. a SQL Server filegroup
- D. an Azure SQL database

Answer: ([SHOW ANSWER](#))

**NEW QUESTION: 24**

Hotspot Question

You have an Azure subscription that contains a SQL Server instance on a virtual machine named Server1.

You execute the following T-SQL statement on Server1.

```

@message_id=0,
@severity=20,
@enabled=1,
@delay_between_responses=0,
@include_event_description_in=1;
)

KEC msdb.dbo.sp_add_operator
@name = 'DBA-Team',
@enabled = 1,
@email_address = 'DBA-Team@contoso.com';
)

KEC msdb.dbo.sp_add_notification
@alert_name = N'Fatal Error',
@operator_name = 'DBA-Team',
@notification_method = 1;

```

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

NOTE: Each correct selection is worth one point.

Answer Area		Microsoft	
Statements	Yes	No	
An alert will be generated when a fatal error occurs.	<input type="radio"/>	<input type="radio"/>	
A SQL Server Agent job will run when an alert is generated.	<input type="radio"/>	<input type="radio"/>	
An email will be sent to DBA-Team@contoso.com when an alert is generated.	<input type="radio"/>	<input type="radio"/>	

Answer:

## Answer Area

### Statements

An alert will be generated when a fatal error occurs.

Yes

No

A SQL Server Agent job will run when an alert is generated.

An email will be sent to DBA-Team@contoso.com when an alert is generated.

### NEW QUESTION: 25

You have an Azure subscription that contains an Azure SQL managed instance named Server1. You need to ensure the availability of Server1 if an Azure region fails.

What should you do?

- A. Migrate the databases hosted on Server1 to the Hyperscale service tier and enable locally-redundant storage (LRS).
- B. Create a failover group and deploy a logical server for Azure SQL Database to a different Azure region.
- C. Create a failover group and deploy a new SQL managed instance to a different Azure region.
- D. Create a read-only replica in a different Azure region and configure Server1 to use the replica.

**Answer: (SHOW ANSWER)**

To protect an Azure SQL Managed Instance from Azure region failures, you should enable zone redundancy for high availability within the region and implement disaster recovery with failover groups to a secondary region. Zone redundancy ensures resilience to zonal outages, while failover groups enable failover to a secondary instance in case of a region-wide outage.

#### 1. Zone Redundancy:

Enable zone redundancy to ensure the managed instance is distributed across multiple availability zones within the same region.

This provides redundancy against hardware and software failures in a single zone, maintaining high availability with up to 99.99% SLA.

You can enable zone redundancy through the Azure portal or PowerShell.

#### 2. Disaster Recovery with Failover Groups:

Create a failover group that includes your primary managed instance and a secondary managed instance in a different region.

The secondary instance serves as a failover target in case of a region-wide outage.

Configure the failover group to replicate data continuously between the instances, ensuring minimal data loss during a failover.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/disaster-recovery-guidance>

### NEW QUESTION: 26

You have an Azure subscription.

You plan to deploy four SQL Server on Azure Virtual Machines instances to two Azure regions. Two instances will be deployed to each region. Each region will host a domain-independent availability group.

You need to recommend a high availability, disaster recovery (HA/DR) solution for the planned deployment. The solution must minimize administrative effort.

What should you include in the recommendation?

- A. a Hyper-V cluster
- B. a workgroup cluster
- C. an Active Directory-detached cluster
- D. a Network Load Balancing (NLB) cluster

**Answer: (SHOW ANSWER)**

SQL Server on Azure Virtual Machines, Configure a workgroup availability group You can create an Active Directory domain-independent cluster with an Always On availability group; this is also known as a workgroup cluster.

Windows Server 2016 introduced a new type of Windows Server failover cluster based on the foundation of the Active Directory-detached cluster: a workgroup cluster. A workgroup cluster allows SQL Server to deploy an availability group on top of a WSFC that doesn't require AD DS. SQL Server requires the use of certificates for endpoint security, just as the database mirroring scenario requires certificates. This type of an availability group is called a domain independent availability group. Deploying an availability group with an underlying workgroup cluster supports the following combinations for the nodes that will make up the WSFC:

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-clusterless-workgroup-configure>

<https://learn.microsoft.com/en-us/sql/database-engine/availability-groups/windows/domain-independent-availability-groups>

### **NEW QUESTION: 27**

You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases. The databases perform slower than expected.

You need to identify whether the performance issue relates to the use of tempdb on sqlsrv1.

What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

**Answer: (SHOW ANSWER)**

The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting.

You can use the Intelligent Insights performance diagnostics log of Azure SQL Database to troubleshoot performance issues.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb-contention>

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-use-diagnostics-log>

## **NEW QUESTION: 28**

### Case Study 3 - Contoso, Ltd 2

#### Overview

Contoso, Ltd. is a clothing retailer based in Seattle. The company has 2,000 retail stores across the United States and an emerging online presence.

The network contains an Active Directory forest named contoso.com. The forest is integrated with an Azure Active Directory (Azure AD) tenant named contoso.com. Contoso has an Azure subscription associated to the contoso.com Azure AD tenant.

#### Existing Environment

##### Transactional Data

Contoso has three years of customer, transaction, operational, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL Server instances contain data from various operations systems. The data is loaded into the instances by using SQL Server Integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time periods. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB. Queries for retail store sales will include the retail store addresses.

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

##### Streaming Twitter Data

The ecommerce department at Contoso develops an Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

#### Planned Changes and Requirements

##### Planned Changes

Contoso plans to implement the following changes:

Load the sales transaction dataset to Azure Synapse Analytics.

Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.

Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

#### Sales Transaction Dataset Requirements

Contoso identifies the following requirements for the sales transaction dataset:

Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.

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

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

Ensure that data storage costs and performance are predictable.

Minimize how long it takes to remove old records.

#### Customer Sentiment Analytics Requirements

Contoso identifies the following requirements for customer sentiment analytics:

Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using row-level security (RLS). The users must be authenticated by using their own Azure AD credentials.

Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.

Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.

Ensure that the data store supports Azure AD-based access control down to the object level.

Minimize administrative effort to maintain the Twitter feed data records.

Purge Twitter feed data records that are older than two years.

#### Data Integration Requirements

Contoso identifies the following requirements for data integration:

Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

Identify a process to ensure that changes to the ingestion and transformation activities can be version-controlled and developed independently by multiple data engineers.

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: (SHOW ANSWER)**

The lifecycle management policy lets you:

- Delete blobs, blob versions, and blob snapshots at the end of their lifecycles Scenario:
- Purge Twitter feed data records that are older than two years.

- Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.
- Minimize administrative effort to maintain the Twitter feed data records.

Incorrect Answers:

A: Time-based retention policy support: Users can set policies to store data for a specified interval. When a time-based retention policy is set, blobs can be created and read, but not modified or deleted. After the retention period has expired, blobs can be deleted but not overwritten.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts>

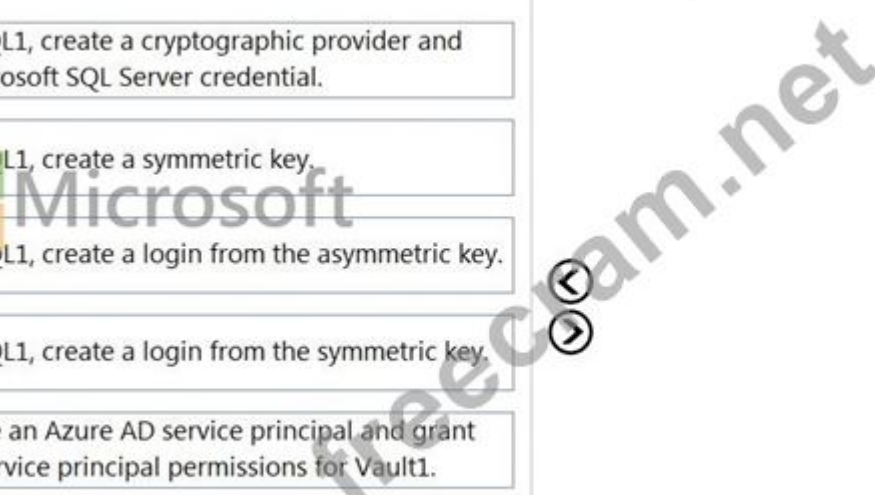
**NEW QUESTION: 29**

Drag and Drop Question

You have an Azure key vault named Vault1 and a SQL server on Azure Virtual Machines instance named SQL1. SQL1 hosts a database named DB1.


You need to configure Transparent Data Encryption (TDE) on DB1 to use a key in Vault1.

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
On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.	<div style="text-align: center;">  </div> <div style="text-align: right; margin-top: 100px;"> <span>⬆</span>  <span>⬇</span>  <span>⬆</span>  <span>⬇</span> </div>
On SQL1, create a symmetric key.	
On SQL1, create a login from the asymmetric key.	
On SQL1, create a login from the symmetric key.	
Create an Azure AD service principal and grant the service principal permissions for Vault1.	
On SQL1, create an asymmetric key.	

**Answer:**

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



**NEW QUESTION: 30**

You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1.

Several users execute adhoc queries to DW1 concurrently.

You regularly perform automated data loads to DW1.

You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run.

What should you do?

- A. Assign a smaller resource class to the automated data load queries.
- B. Create sampled statistics to every column in each table of DW1.
- C. Assign a larger resource class to the automated data load queries.
- D. Hash distribute the large fact tables in DW1 before performing the automated data loads.

**Answer: (SHOW ANSWER)**

The performance capacity of a query is determined by the user's resource class.

Smaller resource classes reduce the maximum memory per query, but increase concurrency.

Larger resource classes increase the maximum memory per query, but reduce concurrency.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workloadmanagement>

**NEW QUESTION: 31**

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

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>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:  
<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**  
**Special Discount Code: freecram**)

### NEW QUESTION: 32

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

SQL1 contains an Extended Events session named session1 that captures Microsoft SQL Server events.

You need to correlate the session events with events captured by Event Tracing for Windows (ETW).

What should you do for session1?

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

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

### NEW QUESTION: 33

## SIMULATION

You need to ensure that any enhancements made to the Query Optimizer through patches are available to db1 and db2 on sql12345678.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

### **Answer:**

Stored procedure execution

You can use the SQL Server Management Studio (SSMS) user interface or Transact-SQL in an SSMS query window to execute a stored procedure. Always use the latest version of SSMS.

Use SQL Server Management Studio

Step 1: In Object Explorer, connect to an instance of SQL Server or Azure SQL Database [here sql12345678], expand that instance, and then expand Databases.

Step 2: Expand the database that you want [Here first db1], expand Programmability, and then expand Stored Procedures.

Step 3: Right-click the stored procedure that you want to run [Here: ALTER DATABASE SCOPED CONFIGURATION] and select Execute Stored Procedure.

Step 4: Add parameter: QUERY\_OPTIMIZER\_HOTFIXES = ON

Step 5: Select OK to execute the stored procedure. If the stored procedure doesn't have any parameters, just select OK.

Step 6: Repeated step 2 to step 5 for db2.

Note: The default approach that SQL Server uses for the query optimizer enables any fixes for the latest database compatibility level for a given product release. This means any fixes for the query optimizer will be used up to the compatibility level of the database, but any hotfixes beyond that compatibility level will not be used. Microsoft provides an option to enable query optimizer hotfixes using the scoped configuration QUERY\_OPTIMIZER\_HOTFIXES.

Use the following command to enable optimizer hotfixes:

```
ALTER DATABASE SCOPED CONFIGURATION QUERY_OPTIMIZER_HOTFIXES = ON;
```

This command configures the database to use all query optimizer hotfixes. For example, optimizations that were included in a recent database cumulative update will apply if the administrator has altered the scope to include query optimizer hotfixes. Using QUERY\_OPTIMIZER\_HOTFIXES applies at the database level, meaning the change would have to be made for each database.

Reference:

<https://infohub.delltechnologies.com/en-US//performance-best-practices-4/sql-server-2019-query-optimizer-hotfixes>

<https://learn.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql>

### **NEW QUESTION: 34**

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

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: 35

### SIMULATION

You need to configure a disaster recovery solution for db1. When a failover occurs, the connection strings to the database must remain the same. The secondary server must be in the West US Azure region.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

Configure active geo-replication and failover (Azure SQL Database)

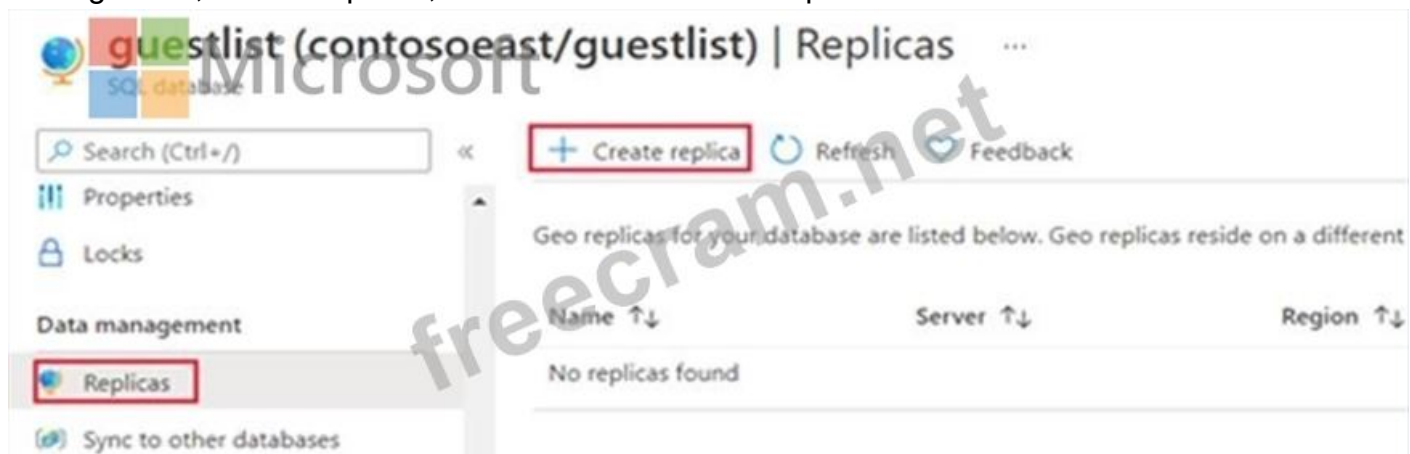
You can configure active geo-replication for Azure SQL Database using the Azure portal or Azure CLI and to initiate failover.

Add a secondary database

The following steps create a new secondary database in a geo-replication partnership.

Step 1: In the Azure portal, browse to the database [Here db1] that you want to set up for geo-replication.

Step 2: On the SQL Database page, select your database [Here db1], scroll to Data management, select Replicas, and then select Create replica.



Step 3: Select or create [Here create] the server for the secondary database, and configure the Compute + storage options if necessary. You can select any region for your secondary server, but we recommend the paired region.

For region select: West US 3

**Create SQL Database - Geo Replica** ...

Microsoft

**Primary database details**

Additional settings will be defaulted where possible based on the the primary database.

Primary database: guestlist

Region: eastus

**Database details**

Enter required settings for this database, including picking a logical server and configuring the compute and storage resources

Database name: guestlist

Server \* ⓘ: contosowest (West US)

Region: West US

Want to use SQL elastic pool? \* ⓘ:  Yes  No

Compute + storage \* ⓘ: **General Purpose**  
Standard-series (Gen5), 2 vCores, 32 GB storage, zone redundant disabled

Step 4: Click Review + create, review the information, and then click Create.

The secondary database is created and the deployment process begins.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-configure-portal>

### NEW QUESTION: 36

Hotspot Question

You have two resource groups named RG1 and RG2. RG1 contains an Azure SQL managed instance named MI1 that hosts a database named DB1. RG2 contains a SQL managed instance named MI2.

You run the following PowerShell script.

```

$database = 'DB1'
$sourceMI = 'MI1'
$sourceRG = 'RG1'
$targetMI = 'MI2'
$targetRG = 'RG2'

## Start database operation.
$databaseCopyParams = @{
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
}
Copy-AzSqlInstanceDatabase @databaseCopyParams

## Verify the operation status is succeeded.

$databaseCopyOperationsParams = @{
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
}
Get-AzSqlInstanceDatabaseCopyOperation @databaseCopyOperationsParams

## Complete database operation.
$completeDatabaseCopyParams = @{
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
}
Complete-AzSqlInstanceDatabaseCopy @completeDatabaseCopyParams

```

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

NOTE: Each correct selection is worth one point.

Answer Area



Statements

The copy of DB1 will be dropped from MI1.

Yes	No
<input type="radio"/>	<input type="radio"/>

A copy of DB1 on MI1 will be available during the operation.

<input type="radio"/>	<input type="radio"/>
-----------------------	-----------------------

A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.

<input type="radio"/>	<input type="radio"/>
-----------------------	-----------------------

Answer:

The screenshot shows the same question as above, but with the correct answers selected. The 'No' radio button for the first statement is selected, the 'Yes' radio button for the second statement is selected, and the 'No' radio button for the third statement is selected.

Explanation:

Box 1: No

The source database DB1 will not be affected.

Box 2: Yes

The source database is available during a Copy-AzSqlInstanceDatabase operation. The operation creates a copy of the source database on a destination instance, and the source database remains online and fully functional throughout the process and after it is completed. The copied database is a transactionally consistent snapshot of the source database at the time the copy request is initiated.

Box 3: No

The Copy-AzSqlInstanceDatabase operation does not utilize backup and restore operations to copy a database. It uses a geo-replication technology to create a transactionally consistent snapshot of the source database.

Reference:

<https://learn.microsoft.com/en-us/powershell/module/az.sql/copy-azsqlinstancedatabase>

**NEW QUESTION: 37**

You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1.

You need to ensure that an automatic email notification is sent once the job completes.

What should you include in the solution?

**A.** From SQL Server Configuration Manager (SSMS), enable SQL Server Agent

**B.** From SQL Server Management Studio (SSMS), run `sp_set_sqlagent_properties`

C. From SQL Server Management Studio (SSMS), create a Database Mail profile

D. From the Azure portal, create an Azure Monitor action group that has an Email/SMS/Push/Voice action

**Answer: (SHOW ANSWER)**

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Using SQL Server Management Studio; to configure SQL Server Agent to use Database Mail:

1. In Object Explorer, expand a SQL Server instance.
2. Right-click SQL Server Agent, and then click Properties.
3. Click Alert System.
4. Select Enable Mail Profile.
5. In the Mail system list, select Database Mail.
6. In the Mail profile list, select a mail profile for Database Mail.
7. Restart SQL Server Agent.

Note: Prerequisites include:

Enable Database Mail.

Create a Database Mail account for the SQL Server Agent service account to use.

Create a Database Mail profile for the SQL Server Agent service account to use and add the user to the DatabaseMailUserRole in the msdb database.

Set the profile as the default profile for the msdb database.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-database-mail>

## NEW QUESTION: 38

Drag and Drop Question

You have an Azure subscription that contains an Azure SQL database named DB1.

You plan to perform a classification scan of DB1 by using Azure Purview.

You need to ensure that you can register DB1.

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 the Purview Studio, create a scan trigger.
- Modify the Access control (IAM) settings for DB1.
- Create a system-assigned managed identity.
- Create a Purview account.
- From the Purview Studio, create a scan rule set.

Answer Area

**Answer:**

Actions

- From the Purview Studio, create a scan trigger.
- From the Purview Studio, create a scan rule set.

Answer Area

- Create a Purview account.
- Modify the Access control (IAM) settings for DB1.
- Create a system-assigned managed identity.

## NEW QUESTION: 39

### Hotspot Question

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.

## Answer Area

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:

<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: 40

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 the on-premises networks shown in the following table.

Name	Location	Subnet
Branch1	Sydney	131.107.10.0/24
Branch2	Melbourne	131.107.11.0/24

You have an Azure subscription that contains an Azure SQL Database server named SQL1.

SQL1 contains two databases named DB1 and DB2.

You need to configure access to DB1 and DB2. The solution must meet the following requirements:

- Ensure that DB1 can be accessed only by users in Branch1.
- Ensure that DB2 can be accessed only by users in Branch2.

Solution: You connect to DB1 and run the following command.

```
EXECUTE sp_set_database_firewall_rule 'Allow db1 users',  
'131.107.10.0', '131.107.10.255'
```

You connect to DB2 and run the following command.

```
EXECUTE sp_set_database_firewall_rule 'Allow db2 users',  
'131.107.11.0', '131.107.11.255'
```

Does this meet the goal?

**A.** Yes

**B.** No

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

#### **NEW QUESTION: 41**

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

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

You have an Azure SQL database named Sales.

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

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

- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and failover groups.

Does this meet the goal?

**A.** Yes

**B.** No

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

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node.

High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

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

## **NEW QUESTION: 42**

### **SIMULATION**

You have a legacy application written for Microsoft SQL Server 2012. The application will be the only application that accesses db1.

You need to ensure that db1 is compatible with all the features and syntax of SQL Server 2012. To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

### **Answer:**

View or change the compatibility level of a database

You can view or change the compatibility level of a database in SQL Server, Azure SQL Database, or Azure SQL Managed Instance by using SQL Server Management Studio or Transact-SQL.

Use SQL Server Management Studio

To view or change the compatibility level of a database using SQL Server Management Studio (SSMS) Step 1: Connect to the appropriate server or instance hosting your database [Here db1].

Step 2: Select the server name in Object Explorer.

Step 3: Expand Databases, and, depending on the database, either select a user database or expand System Databases and select a system database.

Step 4: Right-click the database, and then select Properties.

The Database Properties dialog box opens.

Step 5: In the Select a page pane, select Options.

The current compatibility level is displayed in the Compatibility level list box.

Step 6: To change the compatibility level, select a different option from the list [Select SQL Server 2012]

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/databases/view-or-change-the-compatibility-level-of-a-database>

## **NEW QUESTION: 43**

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

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

#### **NEW QUESTION: 44**

You have an Azure SQL Database elastic pool that contains 10 databases.

You receive the following alert.

Msg 1132, Level 16, state 1, Line 1

The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.

You need to resolve the alert. The solution must minimize administrative effort.

Which three actions can you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

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

<https://learn.microsoft.com/en-us/azure/azure-sql/database/file-space-manage?view=azuresql-db>

#### **NEW QUESTION: 45**

You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1.

You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

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

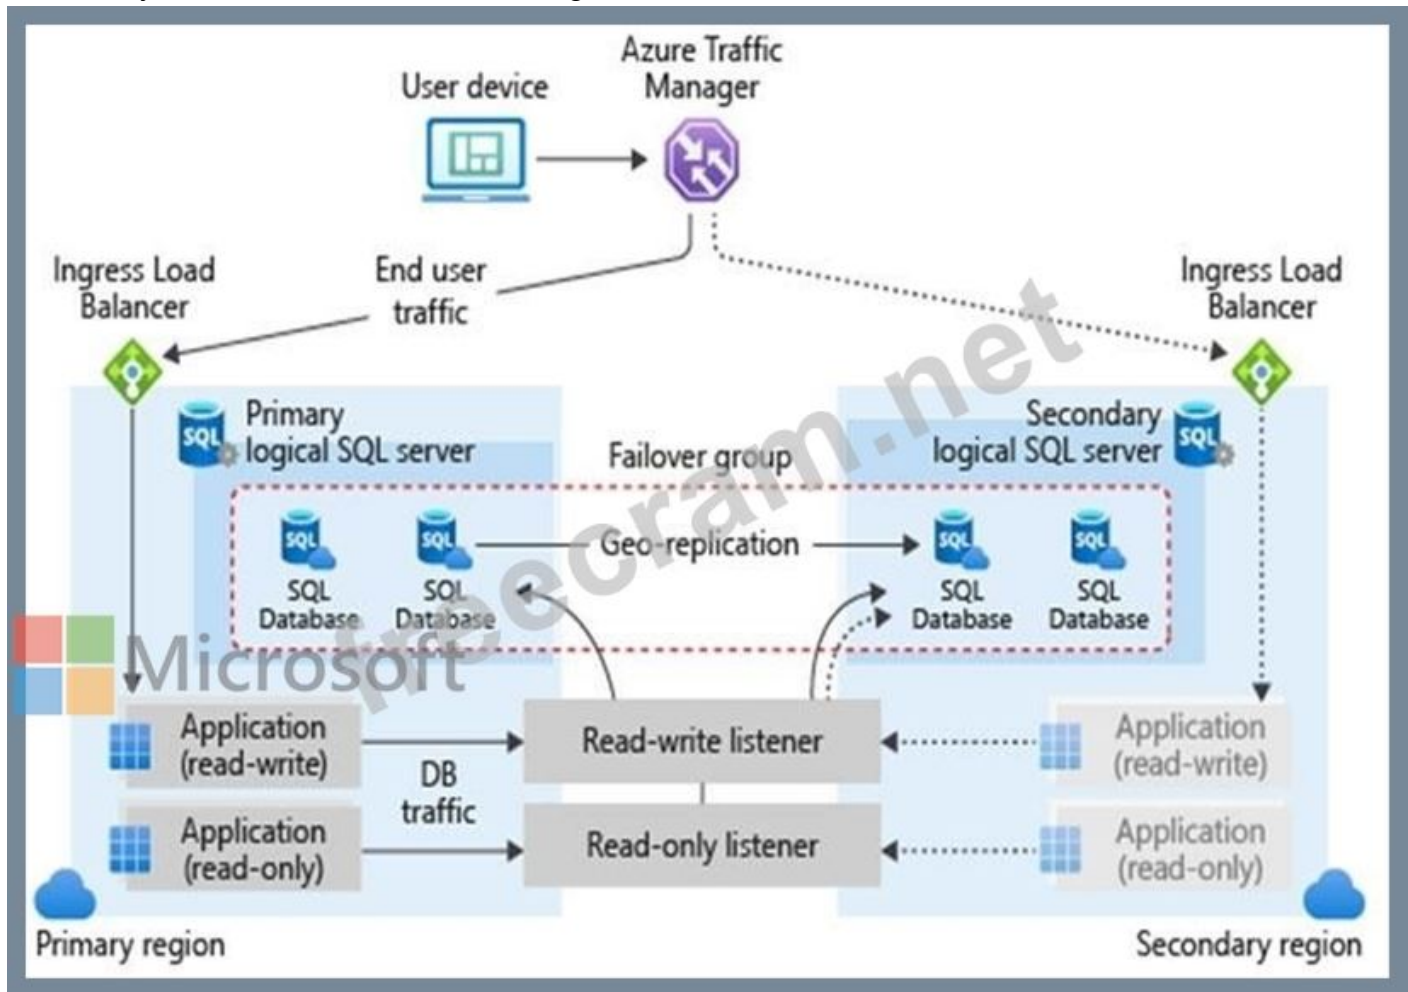
NOTE: Each correct selection is worth one point.

- A. From the properties of DB1. configure geo-replication.
- B. From the properties of Server1 add a failover group.
- C. Create a new Azure SQL Database server named Server2.
- D. From the properties of Server1 configure retention for DB1
- E. Create a new Azure SQL Database instance named DB2.

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

B: The auto-failover groups feature allows you to manage the replication and failover of some or all databases on a logical server to another region.

C: The auto-failover group must be configured on the primary server and will connect it to the secondary server in a different Azure region.



Incorrect:

Not A: Geo-replication failovers require client applications to change their connection strings.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell#best-practices-for-sql-database>

<https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azure-portal>

### NEW QUESTION: 46

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 managed instance General Purpose
- C. Azure SQL Database managed instance Business Critical
- D. Azure SQL Database Hyperscale

Answer: ([SHOW ANSWER](#))

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

Special Discount Code: **freecram**)

### NEW QUESTION: 47

Hotspot Question

You plan to deploy three instances of SQL Server on Azure Virtual Machines that will each contain 20 databases.

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

- Ensures that the deployment is highly available
  - Minimizes administrative effort to manage users, logins, permissions, and SQL Server Agent jobs across the instances
- 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 Area**

Type of availability group:

Always On
Contained
<b>Distributed</b>

Minimum Microsoft SQL Server version:

2017
2019
2022

**Answer:**

## Answer Area



Type of availability group:

- Always On
- Contained
- Distributed

Minimum Microsoft SQL Server version:

- 2017
- 2019
- 2022

### NEW QUESTION: 48

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 the on-premises networks shown in the following table.

Name	Location	Subnet
Branch1	Sydney	131.107.10.0/24
Branch2	Melbourne	131.107.11.0/24

You have an Azure subscription that contains an Azure SQL Database server named SQL1.

SQL1 contains two databases named DB1 and DB2.

You need to configure access to DB1 and DB2. The solution must meet the following requirements:

- Ensure that DB1 can be accessed only by users in Branch1.
- Ensure that DB2 can be accessed only by users in Branch2.

Solution: You connect to DB1 and run the following command.

```
EXECUTE sp_set_firewall_rule 'Deny db1 users', '131.107.11.0', '131.107.11.255'
```

You connect to DB2 and run the following command.

```
EXECUTE sp_set_database_firewall_rule 'Deny db2 users', '131.107.10.0', '131.107.10.255'
```

Does this meet the goal?

A. Yes

B. No

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

This command attempts to deny access from the opposite branch subnets, but it does not explicitly allow the correct subnets access to the respective databases. Moreover, denying a range is not a typical configuration for SQL Database firewall rules, which are generally based on allowing specific IP ranges.

## **NEW QUESTION: 49**

### **SIMULATION**

You need to generate an email alert for db1 if the average CPU percentage utilization is greater than 50 percent for five minutes sampled at one-minute intervals. The alert must be sent to admin@contoso.com.

You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

To create an email alert for high CPU percentage utilization in an Azure database within the Azure portal, you'll first need to navigate to your database resource, then to the "Alerts" section. From there, you'll create a new alert rule, define the CPU percentage threshold, configure an action group for email notifications, and finally, create the alert rule.

Create alerts for Azure SQL Database and Azure Synapse Analytics using the Azure portal Step 1: In the Azure portal, locate the resource you are interested in monitoring and select it.

[Select db1]

Step 2: In the resource menu under Monitoring, select Alerts. The text and icon might vary slightly for different resources.

Step 3: Select the + Create button, then Alert rule.

Step 4: On the Create an alert rule page, the Scope is automatically configured to the individual resource.

Capture activity peaks or sustained resource stress with alerts on the Metrics signal category.

Step 5: On the Condition tab, select See all signals and Select a signal from the list of Metrics.

Step 6: Select the desired metric, for example CPU percentage. Select Apply. [Select CPU

percentage] Step 7: After you select a signal, the Alert logic options appear. A preview of recent activity in this resource for that signal is also displayed.

Step 8: Configure a Threshold type to determine when the alert will take action. Choose Aggregation type, Operator, and Threshold value as desired. A typical threshold is: Static, Maximum, Greater than, 80%.

Threshold: [Select 50%]

Operator: [Select Greater than]

Aggregation granularity: [Select 5 minutes]

Select the interval that's used to group the data points by using the aggregation type function.

Choose an Aggregation granularity (period) that's greater than the Frequency of evaluation to reduce the likelihood of missing the first evaluation period of an added time series.

Frequency of evaluation: [Select one minute]

Select how often the alert rule is to be run. Select a frequency that's smaller than the aggregation granularity to generate a sliding window for the evaluation.

Step 9: Under When to evaluate, determine the desired frequency of evaluation. Use the Check every and Lookback period dropdown lists.

Step 10: Optionally, you can add multiple conditions for this alert, choose the Add condition alert if desired.

Step 11: Select Next: Actions >.

Create the new alert rule

Step 12: Use the Use quick actions feature to create a new action group and provide an email or Azure mobile app notification target.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/alerts-create?>

## **NEW QUESTION: 50**

### **SIMULATION**

You need to configure your user account as the Microsoft Entra admin for the server named sql12345678.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

### **Answer:**

Manage administrator accounts in Azure Active Directory B2C

You can assign a role when you create a user or invite a guest user. You can add a role, change the role, or remove a role for a user:

Solution 1:

Provision Microsoft Entra admin (SQL Managed Instance)

Your SQL Managed Instance needs permission to read Microsoft Entra ID to accomplish tasks such as authentication of users through security group membership or creation of new users. For this to work, you must grant the SQL Managed Instance permission to read Microsoft Entra ID. You can do this using the Azure portal or PowerShell.

To grant your SQL Managed Instance read permissions to Microsoft Entra ID using the Azure portal, sign in as a Global Administrator and follow these steps:

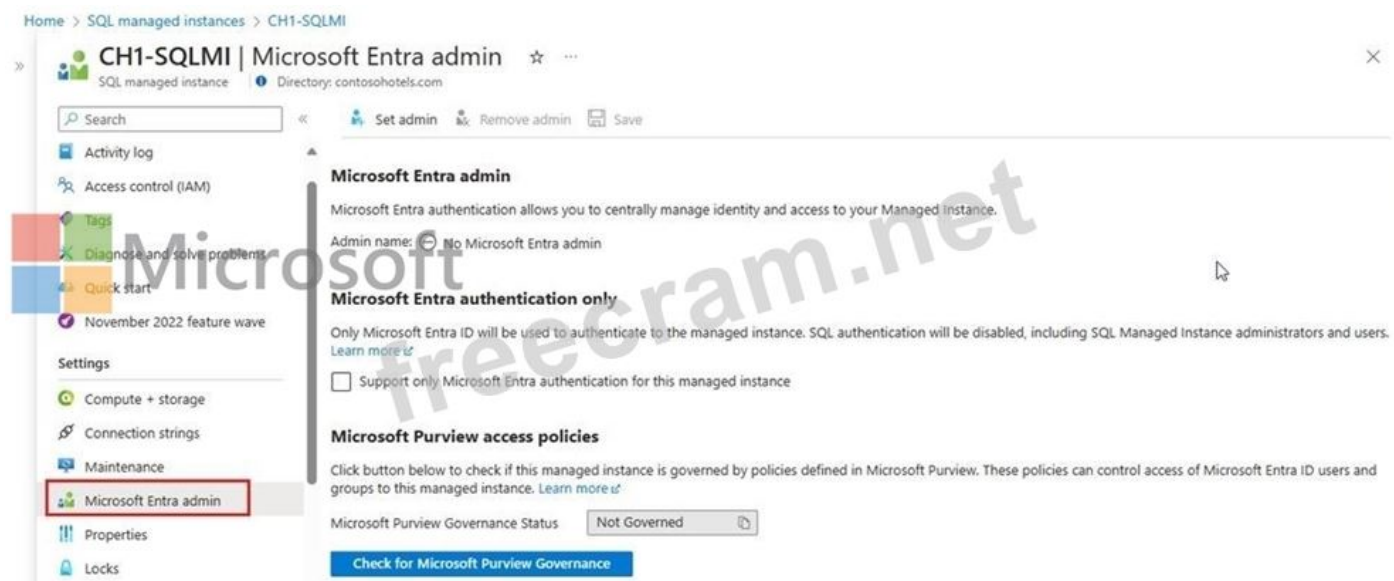
Step 1: In the Azure portal, in the upper-right corner select your account, and then choose Switch directories to confirm which directory is your Current directory. Switch directories, if necessary.



Step 2: Choose the correct Microsoft Entra directory as the Current directory.

This step links the subscription associated with Microsoft Entra ID to the SQL Managed Instance, ensuring the Microsoft Entra tenant and SQL Managed Instance use the same subscription.

Step 3: Now, you can choose your Microsoft Entra admin for your SQL Managed Instance. For that, go to your managed instance resource in the Azure portal and select Microsoft Entra admin under Settings.



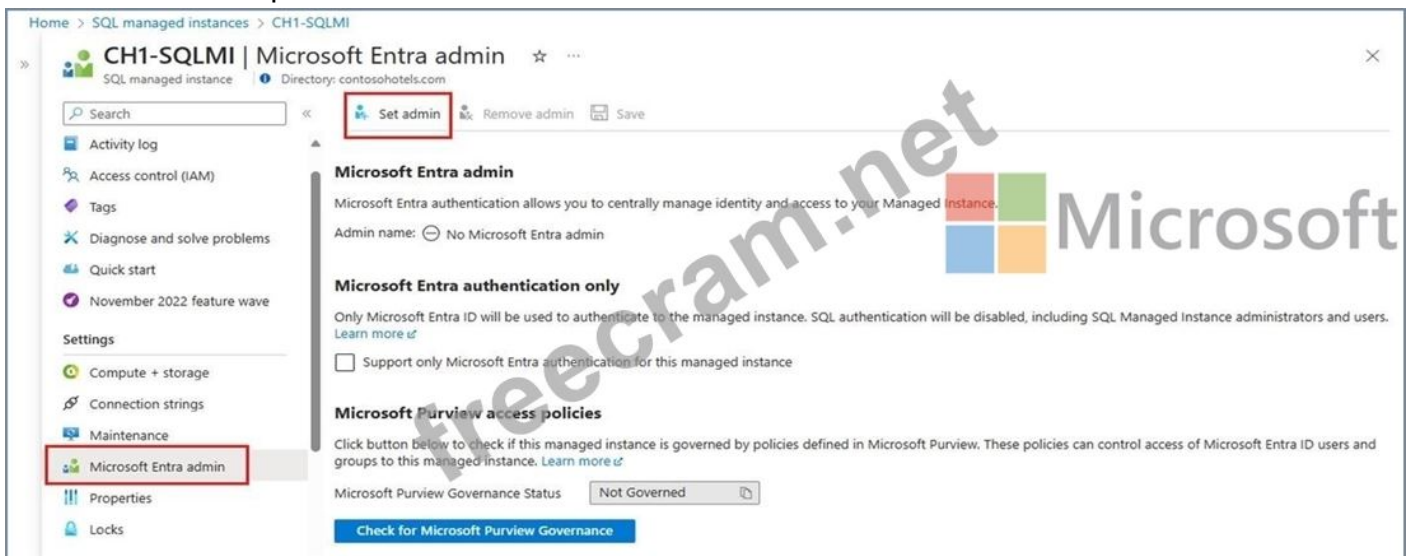
Step 4: Select the banner on top of the Microsoft Entra admin page and grant permission to the current user.



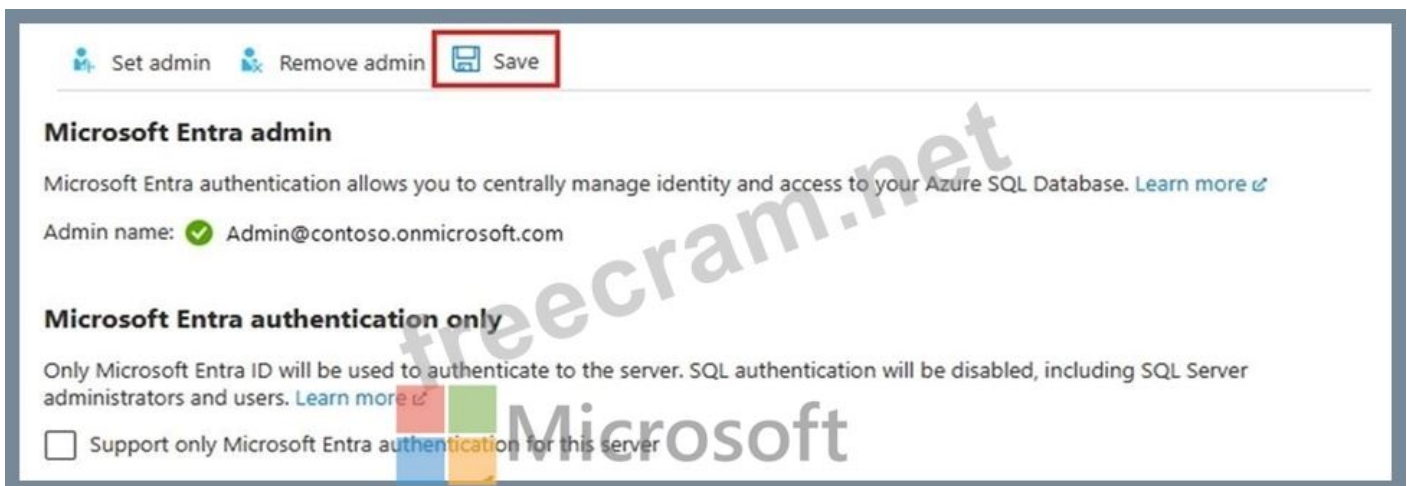
Step 5: After the operation succeeds, the following notification will show up in the top-right corner:



Step 6: On the Microsoft Entra admin page, select Set admin from the navigation bar to open the Microsoft Entra ID pane.



Step 7: On the Microsoft Entra ID pane, search for a user, check the box next to the user or group to be an administrator, and then press Select to close the pane and go back to the Microsoft Entra admin page for your managed instance. [Select yourself] Step 8: From the navigation bar of the Microsoft Entra admin page for your managed instance, select Save to confirm your Microsoft Entra administrator.



The process of changing the administrator might take several minutes. Then the new administrator appears in the Microsoft Entra admin box.

Solution 2:

Step 1: Sign in to the Azure portal with Global Administrator or Privileged Role Administrator permissions.

Step 2: (If you have access to multiple tenants, select the Settings icon in the top menu to switch to your Azure AD B2C tenant from the Directories + subscriptions menu.) Step 3: Under Azure services, select Azure AD B2C. Or use the search box to find and select Azure AD B2C.

Step 4: Under Manage, select Users.

Step 5: Select the user you want to change the roles for. [Select your user account]. Then select Assigned roles.

Step 6: Add a role assignment

Step 7: Select Add assignments, select the role to assign (for example, Application administrator), [Select Microsoft Entra administrator] and then choose Add.

Reference:

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

<https://learn.microsoft.com/en-us/azure/active-directory-b2c/tenant-management-manage-administrator>

### **NEW QUESTION: 51**


Hotspot Question

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  = 'user1@contoso.com'
```

CALLER
LOGIN
OWNER
USER


```
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
```

Explanation:

Box 1: USER

Using EXECUTE AS and REVERT to switch context.

The following example creates a context execution stack using multiple principals. The REVERT statement is then used to reset the execution context to the previous caller.

\*\*

```
EXECUTE AS USER = 'user2';
```

--The following REVERT statements will reset the execution context to the previous context.

```
REVERT;
```

\*\*

Box 2: REVERT

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: 52

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

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

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

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

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

Does this meet the goal?

A. Yes

B. No

**Answer: (SHOW ANSWER)**

You can't overwrite an existing Azure SQL Database database during restore.

Correct solution: From the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups>

### NEW QUESTION: 53

Hotspot Question

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

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

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

NOTE: Each correct selection is worth one point.

Create:   
A materialized view in DB1  
A security policy in the Orders table  
Database scoped credentials in DB1

Control access to the rows by using:   
A masking rule  
A table-valued function  
The CONTAINS predicate

**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:

Implement RLS by using the CREATE SECURITY POLICY Transact-SQL statement, and predicates created as inline table-valued functions.

<https://learn.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver16>

**NEW QUESTION: 54**

### Hotspot Question

You have an Azure subscription. The subscription contains an Azure Storage account named storage1, a container named Container1, and an Azure SQL managed instance that hosts a database named DB1.

You need to back up DB1 to Container1.

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

NOTE: Each correct selection is worth one point.

### Answer Area

BACKUP DATABASE [db1]

TO URL =

- 'abfs://container1.storage1.blob.core.windows.net
- 'https://storage1.blob.core.windows.net/container1
- 'https://storage1.container1.blob.core.windows.net

WITH

- WITH COMPRESSION
- WITH COPY\_ONLY
- WITH DIFFERENTIAL

Answer:

### Answer Area

BACKUP DATABASE [db1]

TO URL =

- 'abfs://container1.storage1.blob.core.windows.net
- 'https://storage1.blob.core.windows.net/container1
- 'https://storage1.container1.blob.core.windows.net

WITH

- WITH COMPRESSION
- WITH COPY\_ONLY
- WITH DIFFERENTIAL

Explanation:

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url>

<https://learn.microsoft.com/en-us/sql/t-sql/statements/backup-transact-sql>

### NEW QUESTION: 55

You have an Azure subscription. The subscription contains an instance of SQL Server on Azure Virtual Machines named SQL1 and an Azure Automation account named account1.

You need to configure account1 to restart the SQL Server Agent service if the service stops. Which setting should you configure?

- A. Start/Stop VM
- B. Change tracking
- C. Update management
- D. State configuration (DSC)

Answer: ([SHOW ANSWER](#))

#### NEW QUESTION: 56

You have an Azure SQL database named DB1.

You need to query the fragmentation information of data and indexes for the tables in DB1.

Which command should you run?

- A. DBCC SHOWCONTIG
- B. sys.dm\_db\_index\_usage\_stats
- C. sts.dm\_db\_index\_physical\_stats
- D. DBCC CHECKALLOC

Answer: ([SHOW ANSWER](#))

#### NEW QUESTION: 57

Drag and Drop Question

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
	From the Azure portal, deploy the templates.
Add IP addresses to the firewall	From the database project, deploy the database schema and permissions

**NEW QUESTION: 58**

Drag and Drop Question

You have an Azure virtual machine named Server1 that contains an instance of Microsoft SQL Server 2022 named SQL1.

SQL1 contains two databases named DB1 and DB2.

You need to take a snapshot backup of DB1 and DB2. The backup must NOT disrupt the backup chain and must NOT affect other databases.

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

**Commands**

```
BACKUP GROUP db1, db2
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
BACKUP SERVER
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2);
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

**Answer Area**



**Answer:**

**Commands**

```
BACKUP SERVER
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

**Answer Area**

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2);
```

```
BACKUP GROUP db1, db2
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF
```



Microsoft

freecram.net

**NEW QUESTION: 59**

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: (SHOW ANSWER)**

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.

Reference:

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

**NEW QUESTION: 60**

You have an Azure virtual machine based on a custom image named VM1.

VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements:

- Automate the patching of SQL Server and Windows Server
- Automate full database backups and transaction log backups of the databases on VM1.
- Minimize administrative effort.

What should you do first?

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

**Answer: (SHOW ANSWER)**

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.

To utilize the SQL IaaS Agent extension, you must first register your subscription with the Microsoft.SqlVirtualMachine provider, which gives the SQL IaaS extension the ability to create resources within that specific subscription.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extension-automate-management> <https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-agent-extension-manually-register-single-vm?tabs=bash%2Cazure-cli>

### NEW QUESTION: 61

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.

You need to ensure that each customer can only connect to and access their respective database.

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

NOTE: Each correct selection is worth one point.

- A. Implement row-level security (RLS).
- B. Create users in each database.
- C. Configure the database firewall.
- D. Configure the server firewall.
- E. Create logins in the master database.
- F. Implement Always Encrypted.

**Answer: (SHOW ANSWER)**

On Azure SQL databases, you do not have to create a Login for the user unless he needs to access several databases within the server. The requirement is each user will only be allowed to access his own database. So with the principle of least privilege, we just need to a contained user. The database firewall is additional security to limit access on each database.

Reference:

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

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

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

<https://www.examdumps.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**)

**Special Discount Code: freecram)**

### NEW QUESTION: 62

Hotspot Question

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

You need to monitor DB1 by using SQL Insights.

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

## Answer Area Microsoft

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:

## Answer Area

To collect monitoring data, use:

	▼
A virtual machine	
An Azure function	
The Azure Monitor agent	

To store monitoring data, create:

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



Explanation:

SQL Insights stores its data in one or more Log Analytics workspaces.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-insights-overview?view=azuresql>

<https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-insights-enable?view=azuresql#create-log-analytics-workspace>

[view=azuresql#create-log-analytics-workspace](https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-insights-enable?view=azuresql#create-log-analytics-workspace)

## NEW QUESTION: 63

### SIMULATION

You need to configure differential backups for the db1 Azure SQL database to be once a day instead of twice a day.

You may need to use SQL Server Management Studio and the Azure portal.

#### Answer:

Change automated backup settings for Azure SQL Database

Change short-term retention policy

You can change the default point-in-time recovery (PITR) backup retention period and the differential backup frequency by using the Azure portal, the Azure CLI, PowerShell, or the REST API.

To change the PITR backup retention period or the differential backup frequency for active databases by using the Azure portal:

Step 1: Go to the logical server in Azure with the databases whose retention period you want to change.

Step 2: Select Backups on the left pane, and then select the Retention policies tab.

Step 3: Select the databases for which you want to change the PITR backup retention. [Select db1]

Step 4: Select Configure policies from the action bar.

Step 5: To change the retention period for point-in-time restore backups, use the slider under Point-in-time restore.

Step 6: To change differential backup frequency, select 12 Hours or 24 hours from the drop-down menu under Differential backup frequency . [Select 12 hours]

The screenshot displays the Azure portal interface for configuring backup policies. The main window shows the 'Backups' section with the 'Retention policies' tab selected. A table lists the databases and their current backup settings:

Database	PITR	Diff. backup freq.
Test_Database_1	7 Days	24 Hours

The 'Configure policies' dialog is open, showing the following settings:

- Point-in-time restore:** A slider is set to 7 days.
- Differential backup frequency:** A dropdown menu is set to 24 Hours.
- Long-term retention:** Weekly LTR Backups are set to 0, and Monthly LTR Backups are set to 0.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/automated-backups-change-settings>

**NEW QUESTION: 64**

You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8.

You need to deploy a failover cluster instance (FCI) to VM1 and VM2 that will use Azure shared disks. The solution must maximize resiliency.

Which quorum option should you use?

- A. node majority with a cloud witness
- B. node majority with no witness
- C. node majority with a file share witness
- D. node majority with a disk witness

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

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/hadr-cluster-best-practices?view=azuresql&tabs=windows2012#quorum> The disk witness is the most resilient quorum option, but to use a disk witness on a SQL Server on Azure VM, you must use an Azure Shared Disk which imposes some limitations to the high availability solution. As such, use a disk witness when you're configuring your failover cluster instance with Azure Shared Disks, otherwise use a cloud witness whenever possible.

**NEW QUESTION: 65****SIMULATION**

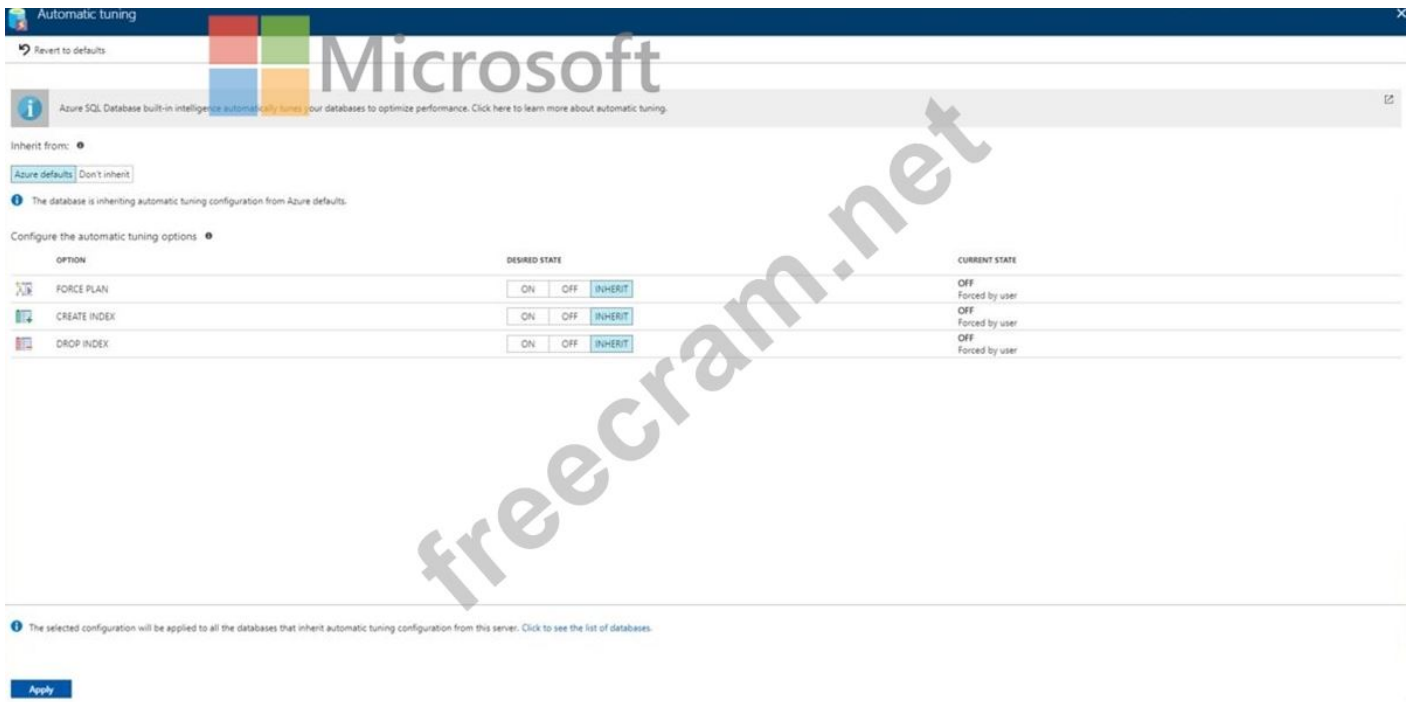
You need to ensure that any unused indexes of the databases on sql12345678 are removed automatically.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

You can go to Azure portal, and enable Azure SQL Database to automatically create or drop indexes:

Step 1: To enable automatic tuning on a server in Azure SQL Database, navigate to the server in the Azure portal. [Here navigate to sql12345678] Step 2: Select Automatic tuning in the menu.



Step 3: Change the setting for "Drop Index" to On.



Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-blog/automatic-index-management-in-azure-sql-database/ba-p/386140>

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

### NEW QUESTION: 66

Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1.

You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance.

Which service tier should you use?

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

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

General purpose has disk latency, Business critical puts the logs on SSD, Premium is a DTU offering which is not supported on Azure SQL MI, Hyperscale is not on Managed Instance.

### NEW QUESTION: 67

Hotspot Question

You have an Azure subscription that is linked to an Azure Active Directory (Azure AD) tenant named contoso.com. The subscription contains an Azure SQL database named SQL1 and an Azure web app named app1. App1 has the managed identity feature enabled.

You need to create a new database user for app1.

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

NOTE: Each correct selection is worth one point.

### Answer Area

The screenshot shows a SQL query editor with the following elements:

- Text: `CREATE USER`
- Dropdown menu for user name: Contains `[App1]`, `[Contoso\app1]`, and `[App1@contoso.com]`.
- Text: `FROM`
- Dropdown menu for authentication type: Contains `login`, `Windows`, and `EXTERNAL PROVIDER`.

**Answer:**

## Answer Area



Explanation:

In T-SQL, just point to App name not managed identity.

<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-sql>

### NEW QUESTION: 68

#### SIMULATION

You need to protect all the databases on sql12345678 from SQL injection attacks.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

#### Answer:

Configure Advanced Threat Protection for Azure SQL Database

Advanced Threat Protection for Azure SQL Database detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Advanced Threat Protection can identify Potential \*SQL injection\*, Access from unusual location or data center, Access from unfamiliar principal or potentially harmful application, and Brute force SQL credentials.

Set up Advanced Threat Protection in the Azure portal

Step 1: Sign into the Azure portal.

Step 2: Navigate to the configuration page of the server you want to protect. [Here sql12345678 ]

In the security settings, select Microsoft Defender for Cloud.

Step 3: On the Microsoft Defender for Cloud configuration page:

Step 3a: If Microsoft Defender for SQL hasn't yet been enabled, select Enable Microsoft Defender for SQL.

Step 3b: Select Configure.

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

Home > chrisqpublictest

chrisqpublictest | Microsoft Defender for Cloud

SQL server

Search (Ctrl+/)

Backups

Deleted databases

Failover groups

Import/Export history

Security

Auditing

Firewalls and virtual networks

Private endpoint connections

Microsoft Defender for Cloud

Transparent data encryption

Identity

Intelligent Performance

Visit Microsoft Defender for Cloud to manage security across your virtual networks, data, apps, and more

Recommendations Security alerts Findings Microsoft Defender for SQL: Disabled (Configure)

5 0 5


Learn more About Microsoft Defender for Cloud About Microsoft Defender for SQL

### Microsoft Defender for SQL

Azure Defender for SQL helps you strengthen your security posture, identify and manage security vulnerabilities and protect against threats on your SQL servers.

You are invited to a 30-day trial, free of charge. After the trial ends, you will be charged \$15/Server/Month

Enable Microsoft Defender for SQL



Step 3c: Under ADVANCED THREAT PROTECTION SETTINGS, select Add your contact details to the subscription's email settings in Defender for Cloud.



Home &gt; chrisqpublictest &gt;

# Server settings

chrisqpublictest



Save



Discard



Feedback

## VULNERABILITY ASSESSMENT SETTINGS

### Subscription

Contoso Team

[Select Subscription](#)

### Storage account

[Select Storage account](#)

Periodic recurring scans

ON

OFF



Send scan reports to ⓘ



Also send email notification to admins and subscription owners ⓘ

## ADVANCED THREAT PROTECTION SETTINGS

Advanced Threat Protection for SQL alerts emails are sent by Defender for Cloud.  
Add your contact details to the subscription's email settings in Defender for Cloud. ⓘ

ⓘ [Enable Auditing for better threats investigation experience](#)

Step 3d: Provide the list of emails to receive notifications upon detection of anomalous database activities in the Additional email addresses (separated by commas) text box. [Skip] Step 3e: Optionally customize the severity of alerts that will trigger notifications to be sent under Notification types. [Skip] Step 3f: Select Save.

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

Home > chrisqpublictest > Server settings >

## Email notifications

Contoso Team

Save

### Email recipients

Select who'll get the email notifications from Defender for Cloud for the Contoso Team subscription.

All users with the following roles:

Additional email addresses (separated by commas):

### Notification types

Use the settings below to select the type of email notifications to be sent by Defender for Cloud.

Notify about alerts with the following severity (or higher):

📘 You'll receive a maximum of one email per 6 hours for high-severity alerts, one email per 12 hours for medium-severity alerts, and one email per 24 hours for low-severity alerts. Learn more >

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/threat-detection-configure>

### NEW QUESTION: 69

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

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

You have an Azure SQL database named Sales.

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

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

- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Does this meet the goal?

A. Yes

B. No

**Answer: (SHOW ANSWER)**

Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different

replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

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

### NEW QUESTION: 70

You have an Azure subscription that contains a Microsoft SQL Server 2022 on Windows Server 2022 virtual machine named Server1. Server1 contains a database named DB1.

You need to back up DB1. The solution must meet the following requirements:

- Minimize how long it takes to perform the backup.
- Minimize the performance impact on Server1.

Which T-SQL statement should you run?

- A. BACKUP LOG db1 TO URL
- B. BACKUP DATABASE db1 WITH FILE\_SNAPSHOT TO URL
- C. BACKUP DATABASE db1 TO URL
- D. BACKUP DATABASE db1 TO DISK

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

A Microsoft SQL Server database log backup copies the transaction log, which records all changes made to the database since the last log backup. This allows you to recover the database to a specific point in time after a full database backup. It's like an incremental backup for SQL Server, preserving only the new changes.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/backup-restore/transaction-log-backups-sql-server>

<https://learn.microsoft.com/en-us/sql/t-sql/statements/backup-transact-sq>

### NEW QUESTION: 71

#### SIMULATION

You plan to perform performance testing of db1.

You need prevent db1 from reverting to the last known good query plan.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

Automatic tuning options

The automatic tuning options available in Azure SQL Database and Azure SQL Managed Instance are:

\* FORCE LAST GOOD PLAN (automatic plan correction)

Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and forces queries to use the last known good plan instead of the regressed plan.

\* Etc.

You can either manually apply tuning recommendations using the Azure portal, or you can let automatic tuning autonomously apply tuning recommendations for you.

Azure defaults are set to FORCE\_LAST\_GOOD\_PLAN enabled.

We need to disable it.

With transact-SQL:

ALTER DATABASE SET options (Transact-SQL)

Parameters include:

<automatic\_tuning\_option> ::=

```
{  
AUTOMATIC_TUNING ( FORCE_LAST_GOOD_PLAN = { DEFAULT | ON | OFF } )  
}
```

Solution with T-SQL

Stored procedure execution

You can use the SQL Server Management Studio (SSMS) user interface or Transact-SQL in an SSMS query window to execute a stored procedure. Always use the latest version of SSMS.

Use SQL Server Management Studio

Step 1: In Object Explorer, connect to an instance of SQL Server or Azure SQL Database, expand that instance, and then expand Databases.

Step 2: Expand the database that you want, expand Programmability, and then expand Stored Procedures.

Step 3: Right-click the stored procedure that you want to run [here: ALTER DATABASE SET] and select Execute Stored Procedure.

Step 4: Set the parameters: AUTOMATIC\_TUNING ( FORCE\_LAST\_GOOD\_PLAN = { OFF } )

Step 5: Select OK to execute the stored procedure. If the stored procedure doesn't have any parameters, just select OK.

Reference:

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

<https://learn.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-set-options>

<https://learn.microsoft.com/en-us/sql/relational-databases/stored-procedures/execute-a-stored-procedure>

## **NEW QUESTION: 72**

Case Study 1 - Litware, Inc

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.

#### Drag and Drop Question

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.

The screenshot shows a Microsoft exam question interface. On the left, under the heading "Actions", there is a list of five actions in rectangular boxes:

- 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.

On the right, under the heading "Answer Area", there is an empty space for the answer. A watermark "freecram.net" is overlaid diagonally across the interface. At the bottom right, the Microsoft logo is visible.

**Answer:**

## Actions

Create an Azure Key Vault instance and generate a secret.

Create an Azure AD managed identity.



## Answer Area

Register ResearchApp1 to Azure AD.

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

Run the Always Encrypted wizard.

Explanation:

Box 1: Register ResearchApp1 to Azure AD.

Box 2: Create an Azure Key Vault instance and configure an access policy.

Need to configure the key vault with an access policy to enable ResearchApp1 retrieval of the keys.

Box 3: Run the Always Encrypt wizard.

Run the Always Encrypt wizard to encrypt the columns and store the encryption keys in the vault.

Reference:

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

## NEW QUESTION: 73

Hotspot Question

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:

Box 1: No

Short term retention is set to 21 days.

Note: Short term retention policy is used to address point-in-time restores whereas a long-term retention policy is used to address restores from long-term or older backups for various audit and compliance purposes. We can also save these backup files as part of a short-term retention policy for up to 7-35 days.

Box 2: Yes

Weekly retention is set to P52W, which stands for a periodic frequency of 52 weeks (364 days).

Note: -WeeklyRetention

The Weekly Retention. If just a number is passed instead of an ISO 8601 string, days will be assumed as the units. There is a minimum of 7 days and a maximum of 10 years.

Box 3: No

Yearly retention is set to PSY.

Note: -YearlyRetention

The Yearly Retention. If just a number is passed instead of an ISO 8601 string, days will be assumed as the units. There is a minimum of 7 days and a maximum of 10 years.

Reference:

<https://strata.opengamma.io/apidocs/com/opengamma/strata/basics/schedule/Frequency.html>

<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: 74**

You have a failover cluster deployed to Microsoft Azure virtual machines.

You plan to deploy a new Microsoft SQL instance to the cluster.

The instance will host one database that has 10 file groups.

You provision 10 disks for the instance.

You need to recommend a storage solution that always provides the fastest possible reads for the database.

What should you include in the recommendation?

- A. dynamic disks that use mirroring
- B. one file group per disk
- C. a storage space
- D. dynamic disks that use striping

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 75**

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

- The web app must be hosted on an Azure virtual network.
- The Azure SQL database must be assigned a private IP address.
- The Azure SQL database must allow connections only from the virtual network.

You need to recommend a solution that meets the requirements.

What should you include in the recommendation?

- A. a database-level firewall
- B. a server-level firewall
- C. Azure Private Link

D. a network security group (NSG)

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

## **NEW QUESTION: 76**

### SIMULATION

You need to encrypt the LastName column of the SalesLT.Customer table in db1 by using Always Encrypted. You must use the Windows Certificate Store.

You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

Solution 1 [Less complex]

Configure column encryption using Always Encrypted Wizard

The Always Encrypted Wizard is a powerful tool that allows you to set the desired Always Encrypted configuration for selected database columns. Depending on the current configuration and the desired target configuration, the wizard can encrypt a column, decrypt it (remove encryption), or re-encrypt it (for example, using a new column encryption key or an encryption type that is different from the current type, configured for the column). Multiple columns can be configured in a single run of the wizard.

Step 1: Open the Always Encrypted Wizard

You can launch the wizard at three different levels:

At a database level - if you want to encrypt multiple columns located in different tables.

At a table level - if you want to encrypt multiple columns located in the same table.

\*-> At a column level - if you want to encrypt one specific column. [This option here] Step 2:

Connect to your SQL Server with the Object Explorer component of SQL Server Management Studio.

Step 3: To encrypt an individual column, navigate to the column, right-click on it, and then select Encrypt Columns. [Select the LastName column of the SalesLT.Customer table in db1] Step 4:

Column Selection page In this page, you select columns you want to encrypt, re-encrypt, or decrypt, and you define the target encryption configuration for the selected columns.

Step 5: If you want the wizard to encrypt or re-encrypt one or more columns using a new column encryption key, pick a key containing (New) in its name. The wizard will generate the key.

Step 6: Master Key Configuration page

If you have selected an autogenerated column encryption key for any column on the previous page, in this page you need to either select an existing column master key or configure a new column master key that will encrypt the column encryption key.

When configuring a new column master key, you can either pick an existing key in Windows Certificate Store or in Azure Key Vault and have the wizard to create just a metadata object for the key in the database, or you can choose to generate both the key and the metadata object describing the key in the database.

Step 7: Click "Next" and follow the wizard's instructions to complete the process.

Solution 2 [A bit more complex. Use for reference]

To encrypt a specific column using Always Encrypted with a Windows Certificate Store, you'll first need to set up a Column Master Key (CMK) [Task 1] and a Column Encryption Key (CEK) [Task 2]. The CMK will be protected by a certificate in the Windows Certificate Store, and the CEK will then be used to encrypt the data in your chosen column.

#### Task 1: Set up a Column Master Key (CMK)

Create a self-signed certificate using SQL Server Management Studio (SSMS) Provision Column Master Keys with the New Column Master Key Dialog The New Column Master Key dialog allows you to generate a column master key or pick an existing key in a key store, and create column master key metadata for the created or selected key in the database.

Step 1: Using Object Explorer, navigate to the Security -> Always Encrypted Keys node under your database.

Step 2: Right-click on the Column Master Keys node and select New Column Master Key....

Step 3: In the New Column Master Key dialog, enter the name of the column master key metadata object.

Step 4: Select a key store:

\*-> Certificate Store - Current User - indicates the Current User certificate store location in the Windows Certificate Store, which is your personal store. [Select this one] Certificate Store - Local computer Azure Key Vault Etc.

Step 5: If you're using SQL Server 2019 (15.x) and your SQL Server instance is configured with a secure enclave, you can select the Allow enclave computations checkbox to make the master key enclave-enabled. [Skip] Step 6: Pick an existing key in your key store, or click the Generate Key or Generate Certificate button, to create a key in the key store. [Pick existing] Step 7: Click OK and the new key appears in the list.

#### Task 2: Create a Column Encryption Key (CEK)

1. In SSMS, navigate to "Always Encrypted Keys" -> "Column Encryption Keys".
2. Right-click and select "New Column Encryption Key".
3. Give the CEK a name.
4. Select the CMK you just created. [Created in Stage 1]
5. Click "OK".

#### Task 3: Encrypt the data in your chosen column.

1. Right-click on the specific column you want to encrypt in Object Explorer. [Select the LastName column of the SalesLT.Customer table in db1]
2. Select "Tasks" -> "Encrypt Columns".
3. The Always Encrypted Wizard will guide you through the process:
4. Select the column you want to encrypt. [Select the LastName column of the SalesLT.Customer table in db1]
5. Choose the encryption type (Deterministic or Randomized).
6. Select the CEK you created. [Created in Stage 2]
7. Click "Next" and follow the wizard's instructions to complete the process.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-wizard>

<https://learn.microsoft.com/en-us/sql/relational-databases/security/encryption/create-and-store-column-master-keys-always-encrypted>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdumps.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

**Special Discount Code: **freecram****)

## NEW QUESTION: 77

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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.

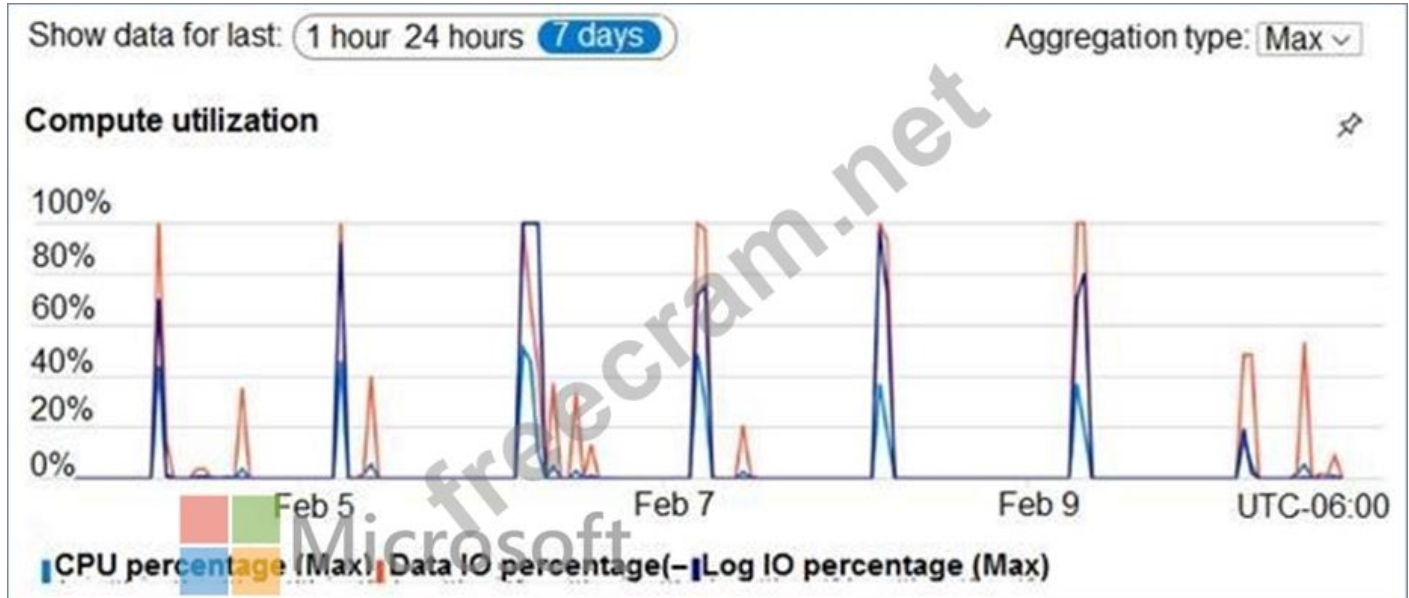
Use a single dashboard to review security and audit data for all the PaaS databases.

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.

The screenshot shows the 'Role assignments' page in the Azure portal. It displays a filtered set of results for 15 total role assignments. Two items are displayed:

Name	Type	Role	Scope
<input type="checkbox"/> DBAGroup1	Group	Contributor	This resource
<input type="checkbox"/> DBAGroup2	Group	SQL DB Contributor	This resource

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_writer` to the new schema.

**D.** For each customer, grant the customer `db_writer` to the existing schema.

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

Scenario: 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.

Members of the `db_ddladmin` fixed database role can run any Data Definition Language (DDL) command in a database.

Incorrect:

Not D: `db_writer` does not have permissions to create database objects.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

## **NEW QUESTION: 78**

### Case Study 3 - Contoso, Ltd 2

#### Overview

Contoso, Ltd. is a clothing retailer based in Seattle. The company has 2,000 retail stores across the United States and an emerging online presence.

The network contains an Active Directory forest named `contoso.com`. The forest is integrated with an Azure Active Directory (Azure AD) tenant named `contoso.com`. Contoso has an Azure subscription associated to the `contoso.com` Azure AD tenant.

#### Existing Environment

##### Transactional Data

Contoso has three years of customer, transaction, operational, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL Server instances contain data from various operations systems. The data is loaded into the instances by using SQL Server Integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time periods. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB. Queries for retail store sales will include the retail store addresses.

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

## Streaming Twitter Data

The ecommerce department at Contoso develops an Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

### Planned Changes and Requirements

#### Planned Changes

Contoso plans to implement the following changes:

Load the sales transaction dataset to Azure Synapse Analytics.

Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.

Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

#### Sales Transaction Dataset Requirements

Contoso identifies the following requirements for the sales transaction dataset:

Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.

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

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

Ensure that data storage costs and performance are predictable.

Minimize how long it takes to remove old records.

#### Customer Sentiment Analytics Requirements

Contoso identifies the following requirements for customer sentiment analytics:

Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using row-level security (RLS). The users must be authenticated by using their own Azure AD credentials.

Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.

Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.

Ensure that the data store supports Azure AD-based access control down to the object level.

Minimize administrative effort to maintain the Twitter feed data records.

Purge Twitter feed data records that are older than two years.

#### Data Integration Requirements

Contoso identifies the following requirements for data integration:

Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

Identify a process to ensure that changes to the ingestion and transformation activities can be version-controlled and developed independently by multiple data engineers.

#### Hotspot Question

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 Area

Table type to store retail store data: ▼

- Hash
- Replicated
- Round-robin

Table type to store promotional data: ▼

- Hash
- Replicated
- Round-robin

**Answer:**

Answer Area

Table type to store retail store data: ▼

- Hash
- Replicated
- Round-robin

Table type to store promotional data: ▼

- Hash
- Replicated
- Round-robin

Explanation:

Box 1: Hash

Scenario:

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

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

Box 2: Round-robin

Scenario:

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads. These are some scenarios where you should choose Round robin distribution:

When you cannot identify a single key to distribute your data.

If your data doesn't frequently join with data from other tables.

When there are no obvious keys to join.

Incorrect Answers:

Replicated: Replicated tables eliminate the need to transfer data across compute nodes by replicating a full copy of the data of the specified table to each compute node. The best candidates for replicated tables are tables with sizes less than 2 GB compressed and small dimension tables.

Reference:

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

### NEW QUESTION: 79

You have an Azure subscription that contains a SQL Server on Linux Azure Virtual Machines instance named VM1. VM1 contains a database named DB1.

You need to create operational performance baselines for DB1 and VM1.

What should you use?

- A. InfluxDB, Collectd, and Grafana
- B. Azure Monitor Log Analytics
- C. Azure Event Hubs
- D. top, grep, and sed

**Answer: (SHOW ANSWER)**

Azure Monitor Logs provides monitoring, alerting, and alert remediation capabilities across cloud and on-premises assets. The Log Analytics virtual machine extension for Linux is published and supported by Microsoft. The extension installs the Log Analytics agent on Azure virtual machines, and enrolls virtual machines into an existing Log Analytics workspace.

### NEW QUESTION: 80

You have an Azure subscription that contains an Azure SQL database named db1.

You need to implement SQL insights for db1.

Which two resources should you create first? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a storage account
- B. a virtual machine
- C. an Azure logic app
- D. an Azure function
- E. a Log Analytics workspace

**Answer: (SHOW ANSWER)**

<https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-insights-enable?view=azuresql>

### NEW QUESTION: 81

Drag and Drop Question

You have an Azure subscription that contains a SQL Server on Azure virtual Machines instance named SQLVM1.

You need to ensure that SQL Server Agent jobs on SQLVM1 can run operating system commands.

Which four T-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
<input type="checkbox"/> USE master	
<input type="checkbox"/> CREATE CREDENTIAL	
<input type="checkbox"/> EXEC dbo.sp_grant_login_to_proxy	
<input type="checkbox"/> EXEC dbo.sp_grant_proxy_to_subsystem	
<input type="checkbox"/> USE msdb	
<input type="checkbox"/> EXEC dbo.sp_add_proxy	

**Answer:**

Statements	Answer Area
<input type="checkbox"/> USE master	<input type="checkbox"/> USE msdb
<input type="checkbox"/> EXEC dbo.sp_grant_login_to_proxy	<input type="checkbox"/> CREATE CREDENTIAL
	<input type="checkbox"/> EXEC dbo.sp_add_proxy
	<input type="checkbox"/> EXEC dbo.sp_grant_proxy_to_subsystem

**NEW QUESTION: 82**

Drag and Drop Question

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:

Actions	Answer Area
	Create a master key on database.
	Create an external data source for Azure Blob storage.
Enable Transparent Data Encryption.	Create an external file format to map the parquet files.
	Create the external table FactSalesOrderDetails.
Load the data to a staging table.	

### Explanation:

To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table.

Step 1: Create a master key on database.

1. Create a master key on the database. The master key is required to encrypt the credential secret.

(Create a database scoped credential for Azure blob storage.)

Step 2: Create an external data source for Azure Blob storage.

2. Create an external data source with CREATE EXTERNAL DATA SOURCE..

Step 3: Create an external file format to map the parquet files.

3. Create an external file format with CREATE EXTERNAL FILE FORMAT.

Step 4. Create an external table FactSalesOrderDetails

4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage>

### NEW QUESTION: 83

Hotspot Question

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.


**Answer Area**

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:**

Answer Area



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

Explanation:

Box 1: ...with CHECKSUM option

Azure Database Migration Service does not initiate any backups, and instead uses existing backups, which you may already have as part of your disaster recovery plan, for the migration. Be sure that you take backups using the WITH CHECKSUM option.

Box 2: An SMB share

For online migrations from SQL Server to SQL Managed Instance using Azure Database Migration Service, you must provide the full database backup and subsequent log backups in the SMB network share that the service can use to migrate your databases.

Reference:

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

## NEW QUESTION: 84

Drag and Drop Question

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.



### Answer:

#### Actions

#### Answer Area

Create a column master key.

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.



### 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.

Incorrect Answers:

Deterministic encryption.

Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns. However, it may also allow unauthorized users to guess information about encrypted values by examining patterns in the encrypted column, especially if there's a small set of possible encrypted values, such as True/False, or North/South/East/West region.

Reference:

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

### **NEW QUESTION: 85**

Hotspot Question

You have an Azure subscription that contains a storage account named databasebackups.

You have an Azure SQL managed instance named DB1.

You need to back up DB1 to databasebackups.

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

NOTE: Each correct selection is worth one point.

## Answer Area

CREATE CREDENTIAL

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

WITH IDENTITY =

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



SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-02T19:30:08Z&spr=https&sv=2021-06-

08&sr=b&sig=B%2FxEYQiOC%4BqyYCeqlMSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'

BACKUP DATABASE DB1

TO URL =

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

WITH

	▼
CHECKSUM	
COMPRESSION	
COPY_ONLY	
DIFFERENTIAL	

Answer:

## Answer Area

CREATE CREDENTIAL

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

WITH IDENTITY =

'DatabaseBackups'
'KeyVault1'
'SHARED ACCESS SIGNATURE'

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

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

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

BACKUP DATABASE DB1

TO URL =

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

WITH

CHECKSUM
COMPRESSION
COPY_ONLY
DIFFERENTIAL

### NEW QUESTION: 86

Hotspot Question

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.

## 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))
```

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.

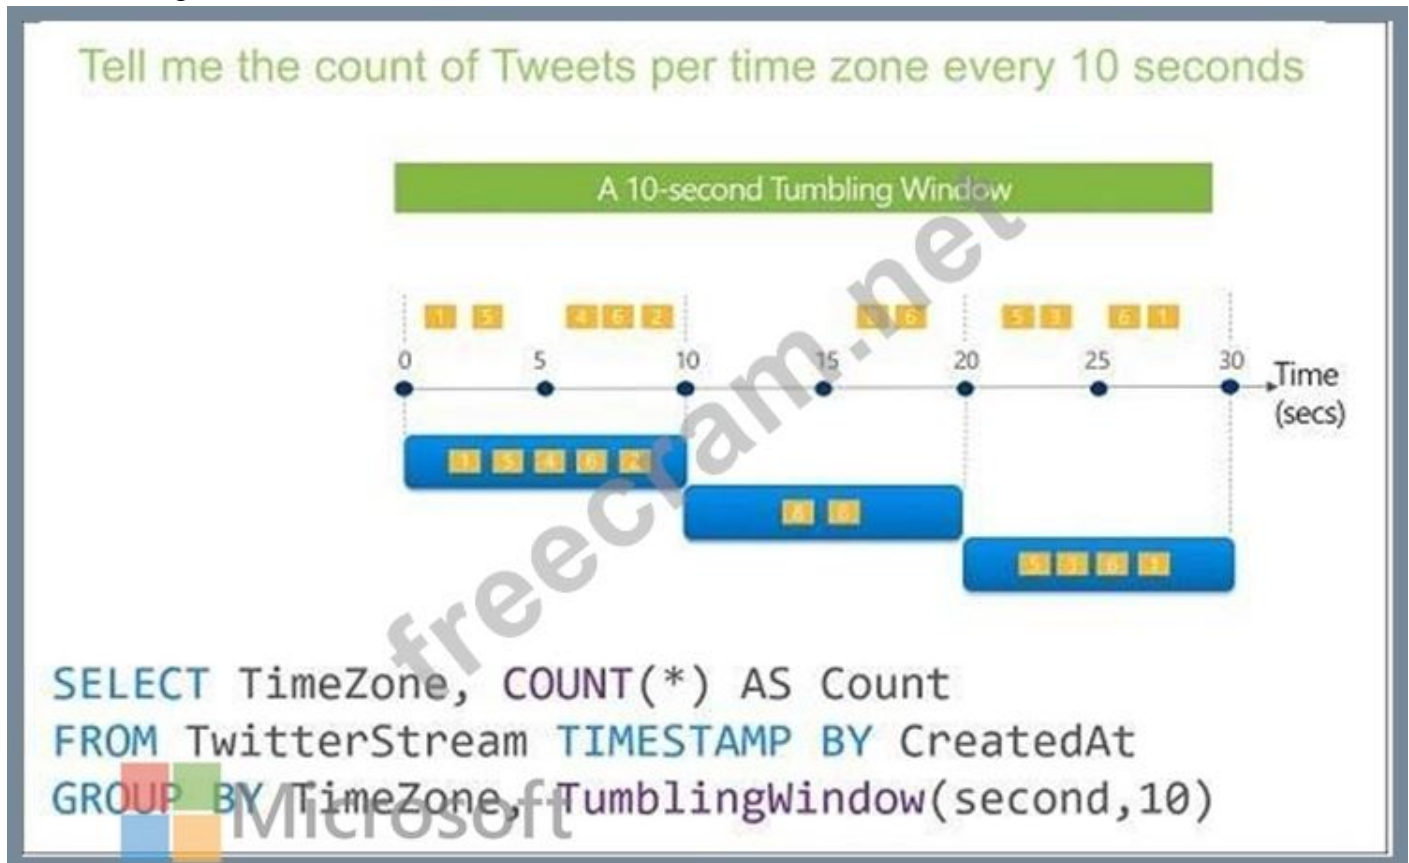
Analytic Function Syntax:

TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT DURATION clause> [<WHEN clause>])

Box 2: Tumbling(minute 5)

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a

Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.



Reference:

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

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

### NEW QUESTION: 87

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-AzSqlDatabase PowerShell cmdlet for Database1 on Server2. You run the Restore-AzSqlDatabase PowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

Correct solution: From the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups>

### NEW QUESTION: 88

Drag and Drop Question

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 in the morning, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00 in the morning.

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

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

**Answer Area**

full backup

Wednesday differential backup

Wednesday log backups

Explanation:

A differential backup is based on the most recent, previous full data backup.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/differential-backups-sql-server?view=sql-server-ver15>

### NEW QUESTION: 89

Case Study 3 - Contoso, Ltd 2

Overview

Contoso, Ltd. is a clothing retailer based in Seattle. The company has 2,000 retail stores across the United States and an emerging online presence.

The network contains an Active Directory forest named contoso.com. The forest is integrated with an Azure Active Directory (Azure AD) tenant named contoso.com. Contoso has an Azure subscription associated to the contoso.com Azure AD tenant.

Existing Environment

Transactional Data

Contoso has three years of customer, transaction, operational, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL Server instances contain data from various operations systems. The data is loaded into the instances by using SQL Server Integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time periods. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB. Queries for retail store sales will include the retail store addresses. You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

#### Streaming Twitter Data

The ecommerce department at Contoso develops an Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

#### Planned Changes and Requirements

##### Planned Changes

Contoso plans to implement the following changes:

Load the sales transaction dataset to Azure Synapse Analytics.

Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.

Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

##### Sales Transaction Dataset Requirements

Contoso identifies the following requirements for the sales transaction dataset:

Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.

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

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

Ensure that data storage costs and performance are predictable.

Minimize how long it takes to remove old records.

##### Customer Sentiment Analytics Requirements

Contoso identifies the following requirements for customer sentiment analytics:

Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using row-level security (RLS). The users must be authenticated by using their own Azure AD credentials.

Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.

Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.

Ensure that the data store supports Azure AD-based access control down to the object level.

Minimize administrative effort to maintain the Twitter feed data records.

Purge Twitter feed data records that are older than two years.

#### Data Integration Requirements

Contoso identifies the following requirements for data integration:

Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

Identify a process to ensure that changes to the ingestion and transformation activities can be version-controlled and developed independently by multiple data engineers.

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: (SHOW ANSWER)**

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-tables-identity>

### NEW QUESTION: 90

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

The output will be sent to a Delta Lake table.

Which output mode should you use?

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

**Answer: (SHOW ANSWER)**

Append Mode: Only new rows appended in the result table since the last trigger are written to external storage. This is applicable only for the queries where existing rows in the Result Table are not expected to change.

Complete Mode: The entire updated result table is written to external storage. It is up to the storage connector to decide how to handle the writing of the entire table.

Update Mode: Only the rows that were updated in the result table since the last trigger are written to external storage. This is different from Complete Mode in that Update Mode outputs only the

rows that have changed since the last trigger. If the query doesn't contain aggregations, it is equivalent to Append mode.

Reference:

<https://docs.databricks.com/getting-started/spark/streaming.html>

## NEW QUESTION: 91

### SIMULATION

You need to enable change data capture (CDC) for db1.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

**Answer:**

Note: Before you can create a capture instance for individual tables, you must enable change data capture for the database.

To enable change data capture, run the stored procedure `sys.sp_cdc_enable_db` (Transact-SQL) in the database context.

#### Stored procedure execution

You can use the SQL Server Management Studio (SSMS) user interface or Transact-SQL in an SSMS query window to execute a stored procedure. Always use the latest version of SSMS.

#### Use SQL Server Management Studio

Step 1: In Object Explorer, connect to an instance of SQL Server or Azure SQL Database, expand that instance, and then expand Databases.

Step 2: Expand the database that you want, expand Programmability, and then expand Stored Procedures.

Step 3: Right-click the stored procedure that you want to run [here: `sys.sp_cdc_enable_db`] and select Execute Stored Procedure.

Step 4: Select OK to execute the stored procedure. If the stored procedure doesn't have any parameters, just select OK.

Note: Enable Database for CDC in Transact-SQL

```
-- =====  
-- Enable Database for CDC  
-- =====  
USE MyDB  
GO  
EXEC sys.sp_cdc_enable_db  
GO
```

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/track-changes/enable-and-disable-change-data-capture-sql-server>

<https://learn.microsoft.com/en-us/sql/relational-databases/stored-procedures/execute-a-stored-procedure>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

Special Discount Code: **freecram**)

### NEW QUESTION: 92

You have an Azure subscription that contains an Azure SQL database. The database contains a table named `table1` that uses partitioned columnstores.

You need to configure `table1` to meet the following requirements:

- Each partition must be compressed.
- The compression ratio must be maximized.
- You must be able to index the compressed data.

What should you use?

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

**Answer: (SHOW ANSWER)**

For columnstore tables and indexes, all columnstore tables and indexes always use columnstore compression and this is not user configurable. 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.

<https://learn.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression?view=sql-server-ver16#:~:text=For%20columnstore%20tables%20and%20indexes%2C%20all,to%20store%20and%20retrieve%20the%20data.>

Both columnstore and columnstore archival compression can be indexed but archival achieves the max compression ratio. There's also no requirement for performance, so the slowness of archival compression is acceptable.

### NEW QUESTION: 93

You have an Azure AD tenant and a logical Microsoft SQL server named `SQL1` that hosts several Azure SQL databases.

You plan to assign Azure AD users permissions to the databases automatically by using Azure Automation.

You need to create the required Automation accounts.

Which two accounts should you create? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. From the Azure Active Directory admin center create a service principal.
- B. From the Azure Active Directory admin center, create a user-assigned managed identity for `SQL1`.
- C. On `SQL1`, create a SQL user in the databases.
- D. On `SQL1`, create a SQL login.
- E. From the Azure Active Directory admin center, create an external identity.

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

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql>

**NEW QUESTION: 94**

You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.

During peak usage, the database will require the following:

- 24 cores
- 500 GB of storage
- 124 GB of memory
- More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard. Which service tier should you use during peak usage?

- A. Business Critical
- B. Premium
- C. Hyperscale

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

In the DTU Purchase Model with "Standard". Business Critical is with General Purpose in vCore Purchase Model.

Migrating a database from the DTU-based purchasing model to the vCore-based purchasing model is similar to scaling between service objectives in the Basic, Standard, and Premium service tiers, with similar duration and a minimal downtime at the end of the migration process. A database migrated to the vCore-based purchasing model can be migrated back to the DTU-based purchasing model at any time in the same fashion, with the exception of databases migrated to the Hyperscale service tier.

<https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-dtu-to-vcore?view=azuresql>

**NEW QUESTION: 95**

Case Study 4 - A.Datum

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.

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

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

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

Scenario:

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

- The sales data including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics.

Streaming Units (SUs) represents the computing resources that are allocated to execute a Stream Analytics job. The higher the number of SUs, the more CPU and memory resources are allocated for your job. This capacity lets you focus on the query logic and abstracts the need to manage the hardware to run your Stream Analytics job in a timely manner.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-streaming-unit-consumption>

### NEW QUESTION: 96

Drag and Drop Question

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.

Values	Answer Area
<input type="text" value="dnsZonePartner"/>	<pre>{   "type": "Microsoft.Sql/managedInstances",   -   },   "identity": {     "type": <input type="text"/>,   },   "dependsOn": [     "[parameters('virtualNetworkName')]"   ],   "properties": {     "administratorLogin": "[parameters('administratorLogin')]",     "administratorLoginPassword": "[parameters('administratorLoginPassword')]",     "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',     'subnetId': "[resourceId('Microsoft.Network/virtualNetworks/subnets',     parameters('virtualNetworkName'), parameters('subnetName'))]", "storageSizeInGB": 8192,     "vCores": 80, "licenseType": "BasePrice",     <input type="text"/> : "True"   } }</pre>
<input type="text" value="storageAccountType"/>	
<input type="text" value="SystemAssigned"/>	
<input type="text" value="UserAssigned"/>	
<input type="text" value="zoneRedundant"/>	



Answer:

**Values****Answer Area**

```

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

```

**NEW QUESTION: 97**

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

You have an Azure subscription that contains an Azure SQL database named SQLDB1.

You need to replicate DB1 to SQLDB1.

Which type of replication should you use?

- A. snapshot
- B. transactional
- C. merge
- D. peer-to-peer

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

**NEW QUESTION: 98**

Hotspot Question

You are working on two projects that will use Azure SQL databases. The projects have the disaster recovery requirements shown in the following table.

Name	Requirement
Project1	<ul style="list-style-type: none"> <li>• Data must sync continuously between primary and secondary replicas.</li> <li>• Read-only secondary replicas must be supported.</li> </ul>
Project2	<ul style="list-style-type: none"> <li>• Data must sync continuously between primary and secondary replicas.</li> <li>• The connection string must remain unchanged after a failover.</li> <li>• Multiple databases must be able to fail over simultaneously.</li> </ul>

You need to recommend a disaster recovery solution for the projects.

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

NOTE: Each correct selection is worth one point.

Answer Area

Project1:

- Active geo-replication
- Always On availability groups
- Failover groups
- Log shipping
- Zone-redundant storage (ZRS)

Project2:

- Active geo-replication
- Always On availability groups
- Failover groups
- Log shipping
- Zone-redundant storage (ZRS)

Answer:

## Answer Area

Project1:

Active geo-replication  
Always On availability groups  
Failover groups  
Log shipping  
Zone-redundant storage (ZRS)

Project2:

Active geo-replication  
Always On availability groups  
Failover groups  
Log shipping  
Zone-redundant storage (ZRS)

Explanation:

Box 1: Active Geo-Replication

Data must sync continuously between primary and secondary replicates.

Read-only secondary replicas must be supported.

For read-only secondary replicas with continuous sync between primary and secondary in Azure SQL Database, you should use Active Geo-Replication. This feature replicates data from a primary database to a readable secondary database, ensuring consistency between them. You can configure this even across different Azure regions for disaster recovery or to scale read workloads.

Box 2: Failover Groups

Data must sync continuously between primary and secondary replicates.

The connection string must remain unchanged after a failover

Multiple databases must be able to fail over simultaneously

To enable multiple Azure SQL databases to fail over simultaneously and sync continuously between primary and secondary replicates, you should use Failover Groups. This solution provides the necessary features for managing the replication and failover of a group of databases.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview>

<https://learn.microsoft.com/en-us/azure/azure-sql/database/failover-group-sql-db>

### NEW QUESTION: 99

You have a new Azure SQL database. The database contains a column that stores confidential information.

You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Turn on auditing and write audit logs to an Azure Storage account.
- B. Add extended properties to the column.
- C. Turn on Advanced Data Security for the Azure SQL server.
- D. Apply sensitivity labels named Highly Confidential to the column.
- E. Turn on Azure Advanced Threat Protection (ATP).

Answer: ([SHOW ANSWER](#))

C: Advanced Data Security (ADS) is a unified package for advanced SQL security capabilities.

ADS is available for Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics. It includes functionality for discovering and classifying sensitive data

D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes that have been added to the SQL Server database engine. This metadata can then be used for advanced, sensitivity-based auditing and protection scenarios.

A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called `data_sensitivity_information`. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	7.125	Microsoft SQL Server Management Studio - Query	1	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview>

### NEW QUESTION: 100

A company plans to use Apache Spark analytics to analyze intrusion detection data.

You need to recommend a solution to analyze network and system activity data for malicious activities and policy violations. The solution must minimize administrative efforts.

What should you recommend?

- A. Azure Data Lake Storage
- B. Azure Databricks

- C. Azure HDInsight
- D. Azure Data Factory

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

Azure DataBricks does have integration with Azure Monitor. Application logs and metrics from Azure Databricks can be send to a Log Analytics workspace.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/databricks-monitoring/application-logs>

**NEW QUESTION: 101**

Hotspot Question

You plan to deploy an instance of SQL Server on Linux Azure Virtual Machines. The instance will run Microsoft SQL Server 2022 and use the SQL Server IaaS Agent extension for Linux.

Which Linux operating system should you deploy, and which benefit will the SQL Server IaaS Agent extension provide? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

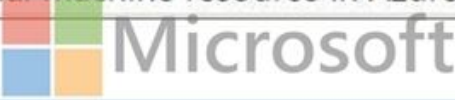
**Answer Area**

Operating system:

- Red Hat Enterprise Linux (RHEL)
- SUSE Linux Enterprise Server (SLES)
- Ubuntu Linux

Benefit:

- Automatic backup
- Automatic patching
- Registration of a SQL virtual machine resource in Azure



**Answer:**

Answer Area

Operating system:

- Red Hat Enterprise Linux (RHEL)
- SUSE Linux Enterprise Server (SLES)
- Ubuntu Linux

Benefit:

- Automatic backup
- Automatic patching
- Registration of a SQL virtual machine resource in Azure



Explanation:

Operating System: Ubuntu Linux

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/linux/sql-server-iaas-agent-extension-linux?view=azuresql&tabs=azure-powershell#limitations> Limitations: The Linux SQL IaaS Agent extension has the following limitations:

Only SQL Server VMs running on the Ubuntu Linux operating system are supported. Other Linux distributions are not currently supported.

Benefit: Registration of a SQL virtual machine resource in Azure

In below link, automated backups and patching are NOT present in the Benefits list for Linux:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/linux/sql-server-iaas-agent-extension-linux?view=azuresql&tabs=azure-powershell>

### **NEW QUESTION: 102**

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

You initiate a copy operation of DB1 to SQLMI2.

You need to monitor the performance of the copy operation.

Which dynamic management view should you use?

- A. sys.dm\_continuous\_copy\_status
- B. sys.dm\_database\_replica\_states
- C. sys.dm\_hadr\_physical\_seeding\_stats
- D. sys.dm\_db\_mirroring\_connections

**Answer: (SHOW ANSWER)**

sys.dm\_hadr\_physical\_seeding\_stats (Transact-SQL)

On the primary replica, query sys.dm\_hadr\_physical\_seeding\_stats DMV to see the physical statistics for each seeding process that is currently running.

Note: Seeding is the process of populating a database with initial data, typically used to provide sample or default records for testing and development purposes. You can use this to create "reproducible environments" for local development, staging, and production.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-database-replica-states-azure-sql-database>

<https://learn.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-hadr-physical-seeding-stats>

### **NEW QUESTION: 103**

You have an Azure subscription that contains a logical SQL server named Server1.

The master database of Server1 contains a user named User1.

You need to ensure that User1 can create databases on Server1.

Which database role should you assign to User1?

- A. db\_owner
- B. dbmanager

C. dbo

D. db\_ddladmin

**Answer: (SHOW ANSWER)**

dbmanager: Can create and delete databases. A member of the dbmanager role that creates a database, becomes the owner of that database, which allows that user to connect to that database as the dbo user. The dbo user has all database permissions in the database. Members of the dbmanager role don't necessarily have permission to access databases that they don't own.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

### NEW QUESTION: 104

Hotspot Question

You have two Azure SQL databases named SQLDB1 and SQLDB2 in the Hyperscale service tier. SQLDB1 was set to the Hyperscale service tier when the database was created. SQLDB2 was changed from the Business Critical service tier to the Hyperscale service tier during the last 14 days.

You need to change SQLDB1 and SQLDB2 to the Business Critical service tier. The solution must minimize the administrative effort.

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

NOTE: Each correct selection is worth one point.

Answer Area



On SQLDB1:  ▼

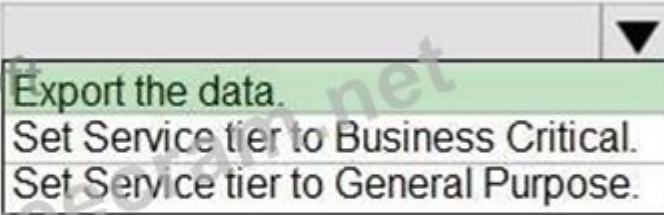
- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

On SQLDB2:  ▼

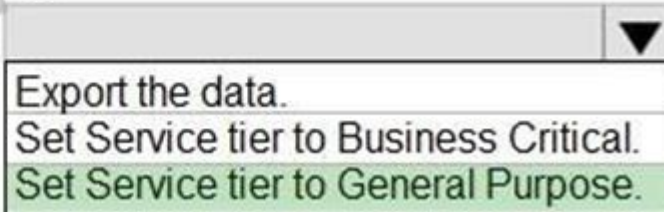
- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

**Answer:**

## Answer Area

On SQLDB1:  ▼

- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

On SQLDB2:  ▼

- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

Explanation:

Box 1: Export the data

For databases that don't qualify for reverse migration [See Limitations for reverse migration below], the only way to migrate from Hyperscale to a non-Hyperscale service tier is to export/import using a bacpac file or other data movement technologies (Bulk Copy, Azure Data Factory, Azure Databricks, SSIS, etc.)

Box 2: Set the Service tier to General Purpose.

Reverse migrate a database from Hyperscale

You can migrate an existing Hyperscale database in Azure SQL Database to the General Purpose service tier using the Azure portal, the Azure CLI, PowerShell, or Transact-SQL. Reverse migration to the General Purpose service tier allows customers who have recently converted an existing database in Azure SQL Database to Hyperscale to move back in an emergency, should Hyperscale not meet their needs. While reverse migration is initiated by a service tier change, it's essentially a size-of-data move between different architectures.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/reverse-migrate-from-hyperscale>

### NEW QUESTION: 105

Hotspot Question

You have an Azure subscription that contains an Azure SQL database named DB1 and a managed identity.

You need to create a user in DB1 that will be used by the managed identity to perform automated tasks.

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

NOTE: Each correct selection is worth one point.

Answer Area



CREATE USER "AutomationAccount"

FROM  ▼  
ASYMETRIC KEY  
EXTERNAL PROVIDER  
FOR CERTIFICATE

WITH  ▼ = '1133a5a2-4c19-47d7-c2d5-42b289997918f'  
KEY  
LOGIN  
OBECT\_ID

EXEC sp\_addrolemember 'db\_owner', "AutomationAccount"

Answer:

Answer Area



CREATE USER "AutomationAccount"

FROM  ▼  
ASYMETRIC KEY  
EXTERNAL PROVIDER  
FOR CERTIFICATE

WITH  ▼ = '1133a5a2-4c19-47d7-c2d5-42b289997918f'  
KEY  
LOGIN  
OBECT\_ID

EXEC sp\_addrolemember 'db\_owner', "AutomationAccount"

Explanation:

Box 1: FROM EXTERNAL PROVIDER

Create a user-assigned managed identity

Box 2: OBJECT\_ID

The above syntax allows creation of Microsoft Entra users without validation. For this to work, the Object Id of the Microsoft Entra principal would have to be supplied and used as an SID in the T-SQL statement.

The validity check of the Object Id is the responsibility of the user running the T-SQL statement.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-azure-ad-user-assigned-managed-identity>

NEW QUESTION: 106

Drag and Drop Question

You have an Azure SQL database named DB1.

You need to create a partitioned table in DB1.


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.

**Objects**

- a partition scheme
- an aligned index
- a filegroup
- a table
- a partition function

**Answer Area**

Navigation icons: left arrow, right arrow, up arrow, down arrow.



**Answer:**

**Objects**

- an aligned index
- a filegroup

**Answer Area**

- a partition function
- a partition scheme
- a table

Navigation icons: left arrow, right arrow, up arrow, down arrow.



Explanation:

<https://learn.microsoft.com/en-us/sql/relational-databases/partitions/partitioned-tables-and-indexes?view=sql-server-ver16>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

**Special Discount Code: freecram**)

**NEW QUESTION: 107**

You have an Azure subscription. The subscription contains three virtual machines that run Red Hat Enterprise Linux (RHEL).

You plan to implement a highly available deployment of Microsoft SQL Server 2022 on the virtual machines by using failover cluster instances (FCIs).

You need to prepare for the implementation.

What should you do first on each server?

- A. Create the FCI disk and network resources.
- B. Install a named SQL Server instance.
- C. Install the cluster resource manager and the FCI resource agent.
- D. Install the default SQL Server instance.

**Answer: (SHOW ANSWER)**

To prepare for implementing a highly available deployment of Microsoft SQL Server 2022 on the virtual machines using failover cluster instances (FCIs), the first step you should take on each server is install the cluster resource manager and the FCI resource agent. Before setting up a failover cluster instance, you need to ensure that the cluster infrastructure is in place. This involves installing the cluster resource manager (such as Pacemaker on Linux) and the FCI resource agent, which are essential for managing the failover clustering and ensuring high availability.

### **NEW QUESTION: 108**

Hotspot Question

You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that contains two databases named DB1 and DB2. The database and log files for DB1 and DB2 are hosted on managed disks.

You need to perform a snapshot backup of DB1 and DB2.

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

NOTE: Each correct selection is worth one point.

## Answer Area

ALTER

DATABASE
SERVER
WORKLOAD GROUP

CONFIGURATION

SET SUSPEND\_FOR\_SNAPSHOT\_BACKUP = ON

(GROUP = (db1, db2));

BACKUP GROUP db1, db2

TO DISK = 'f:\backup\db.bkm'

WITH

COPY_ONLY
FILE_SNAPSHOT
METADATA_ONLY

, FORMAT;



Answer:

**Answer Area**

ALTER  CONFIGURATION

DATABASE  
SERVER  
WORKLOAD GROUP

SET SUSPEND\_FOR\_SNAPSHOT\_BACKUP = ON  
(GROUP = (db1, db2));  
BACKUP GROUP db1, db2  
TO DISK = 'f:\backup\db.bkm'

Microsoft WITH  , FORMAT;

COPY\_ONLY  
FILE\_SNAPSHOT  
METADATA\_ONLY

**NEW QUESTION: 109**

You have four Azure virtual machines hosted in different availability zones in the same Azure region.

You plan to deploy a Microsoft SQL Server Always On failover cluster instance (FCI) to the four virtual machines.

What is the most resilient quorum option for the deployment?

- A. a cloud witness
- B. Node and Disk Majority
- C. Node and File Share Majority
- D. Node Majority

**Answer: (SHOW ANSWER)**

A quorum mode is configured at the WSFC cluster level that dictates the methodology used for quorum voting. The Failover Cluster Manager utility will recommend a quorum mode based on the number of nodes in the cluster.

When using an asymmetric storage configuration for Always On availability groups, you should generally use the Node Majority quorum mode when you have an odd number of voting nodes, or the Node and File Share Majority quorum mode when you have an even number of voting nodes.

Reference:

<https://learn.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/wsfcr-quorum-modes-and-voting-configuration-sql-server>

### NEW QUESTION: 110

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)**

Column-level security simplifies the design and coding of security in your application, allowing you to restrict column access to protect sensitive data.

Reference:

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

### NEW QUESTION: 111

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 the on-premises networks shown in the following table.

Name	Location	Subnet
Branch1	Sydney	131.107.10.0/24
Branch2	Melbourne	131.107.11.0/24

You have an Azure subscription that contains an Azure SQL Database server named SQL1. SQL1 contains two databases named DB1 and DB2.

You need to configure access to DB1 and DB2. The solution must meet the following requirements:

- Ensure that DB1 can be accessed only by users in Branch1.
- Ensure that DB2 can be accessed only by users in Branch2.

Solution: You connect to the master of SQL1 and run the following command.

```
EXECUTE sp_set_firewall_rule 'Allow db1 and db2 users', '131.107.11.0', '131.107.11.255'
```

Does this meet the goal?

- A. Yes

B. No

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

This command sets a firewall rule at the server level (master database), allowing both DB1 and DB2 to be accessed from the Branch2 subnet. This does not meet the requirement of restricting DB1 access to Branch1 only.

### NEW QUESTION: 112

You have an Azure subscription.

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

- Minimize how long it takes to fail over.
- Maintain existing connections to the primary replica during a failover.

What should you do?

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

**Answer: (SHOW ANSWER)**

To match the on-premises experience for connecting to your availability group listener, deploy your SQL Server VMs to multiple subnets within the same virtual network. Having multiple subnets negates the need for the extra dependency on an Azure Load Balancer, or a distributed network name (DNN) to route your traffic to your listener. ref: <https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-overview?view=azuresql-vm#connectivity>

### NEW QUESTION: 113

Hotspot Question

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.

# Answer Area

Encryption key:

	▼
Column encryption key	
Database encryption key	
Service master key	

Encryption type:

	Microsoft ▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	

Answer:

<b>Answer Area</b>	
Encryption key:	▼
Column encryption key	
Database encryption key	
Service master key	
Encryption type:	▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	

Explanation:

Box 1: Column encryption Key

Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Incorrect Answers:

TDE encrypts the storage of an entire database by using a symmetric key called the Database Encryption Key (DEK).

Box 2: Deterministic

Always Encrypted is a feature designed to protect sensitive data, such as credit card numbers or national identification numbers (for example, U.S. social security numbers), stored in Azure SQL Database or SQL Server databases. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to the Database Engine (SQL Database or SQL Server).

Always Encrypted supports two types of encryption: randomized encryption and deterministic encryption. Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.

Incorrect Answers:

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.

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application.

Reference:

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

### **NEW QUESTION: 114**

You have an Azure subscription that contains an instance of SQL Server on an Azure virtual machine named VM1 and an Azure Active Directory Domain Services (Azure AD DS) domain that contains two users named User1 and User 2.

On the default instance of SQL Server on VM1, you create a credential named Credential1 for User1.

You need to ensure that User2 can create a SQL Server Agent proxy that will use Credential1. The solution must use the principle of least privilege.

Which role should you assign to User2?

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

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 115**

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:  $\text{MAX}(\text{<Total number of DBs X average vCore utilization per DB>}, \text{<Number of concurrently peaking DBs X Peak vCore utilization per DB>})$  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: 116**

You have an Azure subscription that contains a Basic SKU Azure SQL database named DB1. You need to configure the long-term retention of backups for DB1. The solution must maximize how long the backups are retained.

Which retention period should you configure?

- A. 180 days
- B. 365 days
- C. five years
- D. 10 years

**Answer: (SHOW ANSWER)**

Long-term retention - Azure SQL Database and Azure SQL Managed Instance Long-term retention can be configured for up to 10 years on backups for Azure SQL Database (including in the Hyperscale service tier), and Azure SQL Managed Instance.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/long-term-retention-overview?view=azuresql>

### NEW QUESTION: 117

Your on-premises network contains a server that hosts a 60-TB database named DB1. The network has a 10-Mbps internet connection.

You need to migrate DB1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

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

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

Database bigger than 40TB, better to use Azure Data Box.

<https://learn.microsoft.com/en-us/azure/databox/data-box-overview>

### NEW QUESTION: 118

Hotspot Question

You have an Azure subscription that contains an Azure SQL database named DB1.

You execute the following T-SQL statements against DB1.

```
CREATE TABLE table1 (  
    column1 TINYINT,  
    column2 VARCHAR(255),  
    column3 NCHAR(10),  
    column4 BINARY(50),  
    column5 DATETIME2  
);
```

```
ALTER TABLE table1 REBUILD WITH (DATA_COMPRESSION = ROW);
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

All the fields in column3 will occupy 10 bytes on disk.



Yes

No

Dictionary compression will be used to reduce the size of table1 on disk.

The data types stored in column1, column2, and column4 support compression.

Answer:

Answer Area

Statements

All the fields in column3 will occupy 10 bytes on disk.

Yes

No

Dictionary compression will be used to reduce the size of table1 on disk.

The data types stored in column1, column2, and column4 support compression.

Explanation:

The NCHAR(10) data type is a fixed-length Unicode character data type. Each character in NCHAR takes 2 bytes because it uses UTF-16 encoding. Therefore, NCHAR(10) will always occupy 20 bytes on disk (2 bytes per character for 10 characters), not 10 bytes. However, the statement is asking if each field occupies 10 bytes, which is misleading; it should be 20 bytes. The statement "ALTER TABLE table1 REBUILD WITH (DATA\_COMPRESSION = ROW);" specifies row-level compression, not dictionary compression. Row compression reduces the size of the data by storing fixed-length data types in variable-length storage format and by reducing metadata overhead. Dictionary compression is a feature of page compression, not row compression.

Column1 (TINYINT): Supports compression. Row compression can store this as a smaller size if possible. Column2 (VARCHAR(255)): Supports compression. Row compression can store only the used length of the string. Column4 (BINARY(50)): Supports compression. Row compression can store binary data more efficiently by removing unnecessary padding.

**NEW QUESTION: 119**

You have an Azure SQL Database managed instance.  
The instance starts experiencing performance issues.

You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.

What should you use?

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

**Answer: (SHOW ANSWER)**

Use the dynamic management view `sys.dm_exec_requests` to track currently executing queries and the associated worker time.

Incorrect:

Not C: DMVs that track Query Store and wait statistics show results for only successfully completed and timed-out queries.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/identify-query-performance-issues>

## **NEW QUESTION: 120**

Hotspot Question

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 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")

Microsoft

Answer:

## 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")

Microsoft

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: 121

You have an Azure SQL managed instance named SQLMI1 that has the following settings:

- vCores: 4
- Service tier: General Purpose
- Hardware generation: Standard-series (Gen5)

You discover that memory pressure on SQLMI1 is high.

You need to reduce the memory pressure on SQLMI1. The solution must minimize costs.

What should to do?

- A. Enable the Query Store.
- B. Change vCores to 8.
- C. Change Hardware generation to Premium-series.
- D. Change Service tier to Business Critical.

**Answer: (SHOW ANSWER)**

The Query Store option is a cost-effective way to reduce memory pressure on a managed Azure SQL instance. The Query Store stores temporary data used by queries in the cloud, reducing the memory requirements of the managed Azure SQL instance.

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

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

<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**)

**Special Discount Code: freecram)**

### NEW QUESTION: 122

You have an application that uses a Microsoft SQL Server database on a Microsoft Azure virtual machine.

The application experiences performance issues, which you suspect are related to the connection pooling.

The issues are prevalent only when there are more than 150 concurrent connections. You need to identify whether the performance issues are related to the connection pooling. Which dynamic management view should you query?

- A. sys.dm\_exec\_requests
- B. sys.dm\_os\_memory\_pools
- C. sys.dm\_exec\_connections
- D. sys.dm\_exec\_sessions

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

### **NEW QUESTION: 123**

You have an on-premises Microsoft SQL Server 2022 instance named SQL1. You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1. You need to configure a bidirectional disaster recovery solution between SQL1 and SQLMI1. What should you include in the solution?

- A. Azure Database Migration Service
- B. Log Replay Service (LRS)
- C. a failover group
- D. the Managed Instance link

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

Disaster recovery with Managed Instance link - Azure SQL Managed Instance You can configure a hybrid disaster recovery solution between SQL Server hosted anywhere and Azure SQL Managed Instance by using the Managed Instance link.

With SQL Server 2022, either SQL Server or Azure SQL Managed Instance can be the primary and you can establish the link initially from either SQL Server or SQL Managed Instance. You can fail over between SQL Server and Azure SQL Managed Instance in either direction, as needed.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/managed-instance-link-disaster-recovery>

### **NEW QUESTION: 124**

You have an Always On availability group deployed to Azure virtual machines. The availability group contains a database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica.

You need to initiate a full backup of DB1 on SQL2.

Which statement should you run?

- A. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (Differential, STATS=5, COMPRESSION);`
- B. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (COPY_ONLY, STATS=5, COMPRESSION);`
- C. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (File_Snapshot, STATS=5, COMPRESSION);`

**D.** BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (NoInit, STATS=5, COMPRESSION);

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

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: 125**

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

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

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

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

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

Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column.

Does this meet the goal?

**A.** Yes

**B.** No

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

Instead, in an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Note: You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the Get Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity>

### **NEW QUESTION: 126**

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

Solution: You change the data file for the master database to autogrow by 10 percent.

Does this meet the goal?

A. Yes

B. No

**Answer: (SHOW ANSWER)**

You should instead reduce the use of table variables and temporary tables.

Or you could create additional tempdb files

Note: The following operations use tempdb extensively:

- \* Repetitive create-and-drop operation of temporary tables (local or global).
- \* Table variables that use tempdb for storage.
- \* Etc.

Reference:

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

### **NEW QUESTION: 127**

You plan to deploy an Azure SQL managed instance.

You need to restore database backups across regions.

Which type of storage account should you use?

A. zone-redundant storage (ZRS)

B. geo-zone-redundant storage (GZRS)

C. locally-redundant storage (LRS)

D. geo-redundant storage (GRS)

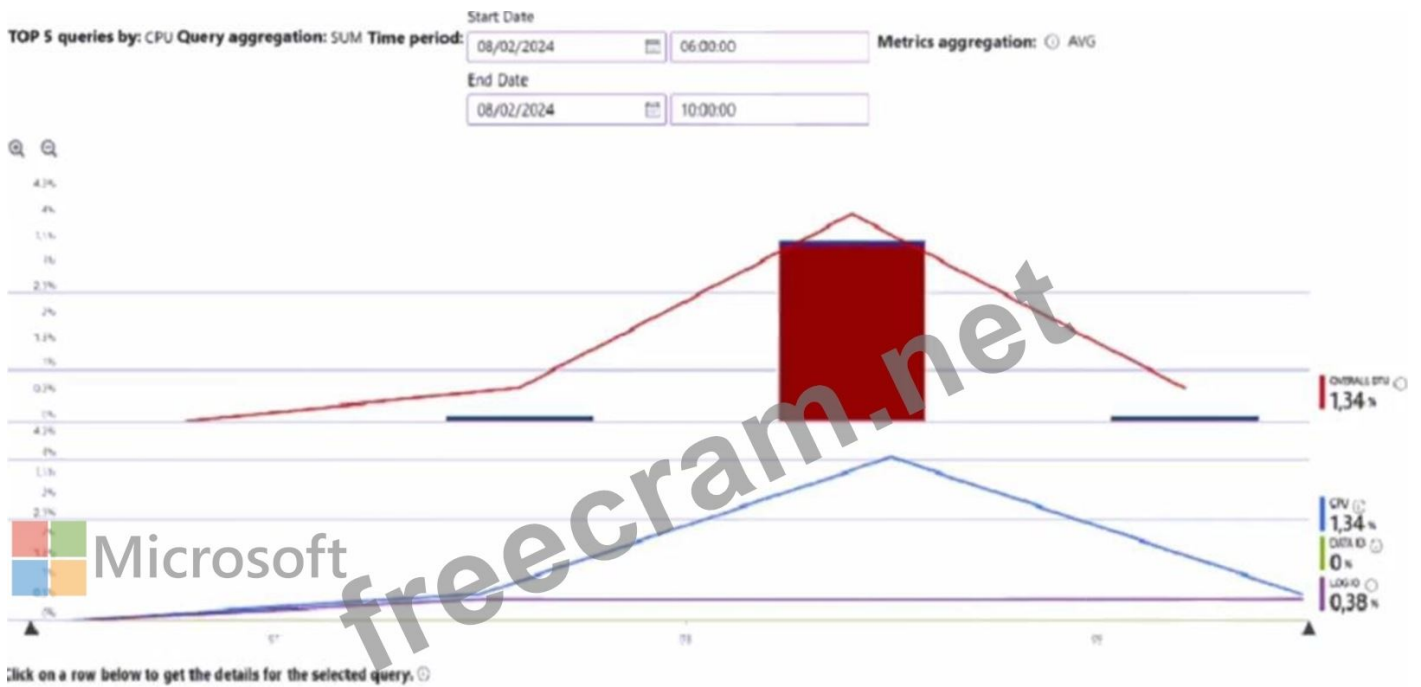
**Answer: (SHOW ANSWER)**

### **NEW QUESTION: 128**

Hotspot Question

You have an Azure subscription that contains an Azure SQL database named DB1.

You have the Query Performance Insight analysis shown in the following exhibit.



Click on a row below to get the details for the selected query.

QUERY ID	CPU[%]	DATA IO[%]	LOG IO[%]	DURATION(hh:mm:ss)	EXECUTIONS COUNT
3964	0.86	0	0	00:04:00,140	8
21195	0.03	0	0.02	00:00:01,180	12713
21196	0.02	0	0.01	00:00:00,850	9297
21196	0.01	0	0	00:00:00,250	2511
21197	0.01	0	0	00:00:00,630	19251

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

Answer Area

Statements

The peak DTU consumption is 4 percent.

Yes No

The peak consumption of the query that has ID 3964 is 4 percent.

The query that has ID 21195 has a duration of 00:00:01.180 every time it runs.

Answer:

Answer Area

Statements

The peak DTU consumption is 4 percent.

Yes No

The peak consumption of the query that has ID 3964 is 4 percent.

The query that has ID 21195 has a duration of 00:00:01.180 every time it runs.

Explanation:

Box 1: No

In the exhibit we see: TOP 5 queries by CPU, Metrics aggregation: AVG

The graphic display average query CPU usage, not the peak DTU consumption.

Box 2: Yes

This is the red line.

Box 3: No

It is the average runtime.

Note: The bottom grid shows aggregated information for the visible queries:

Query ID, which is a unique identifier for the query in the database.

CPU per query during an observable interval, which depends on the aggregation function.

\*-> Duration per query, which also depends on the aggregation function.

Total number of executions for a specific query.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/query-performance-insight-use>

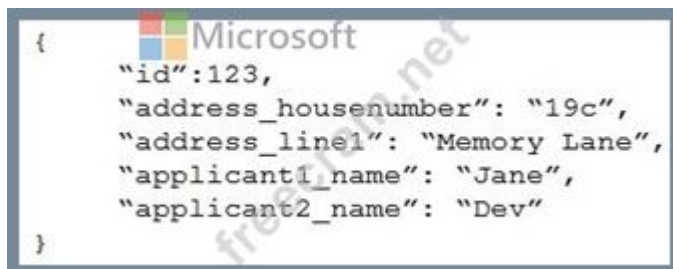
## NEW QUESTION: 129

Hotspot Question

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.

## Answer Area



▼ 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',  
CROSS APPLY  
OPENJSON  
OPENROWSET

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

GO

Answer:

## Answer Area

▼ 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',  
CROSS APPLY  
OPENJSON  
OPENROWSET

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

GO

Explanation:

#### Box 1: CREATE EXTERNAL TABLE

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

Syntax:

```
CREATE EXTERNAL TABLE { database_name.schema_name.table_name |  
schema_name.table_name | table_name }  
( <column_definition> [ ,...n ] )  
WITH (  
LOCATION = 'folder_or_filepath',  
DATA_SOURCE = external_data_source_name,  
FILE_FORMAT = external_file_format_name
```

#### Box 2. OPENROWSET

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

Example:

AS

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

```
'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/  
release/us_population_county/year=*/*.parquet', FORMAT='PARQUET') AS [r] GROUP BY  
decennialTime, stateName GO Reference:
```

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

### **NEW QUESTION: 130**

Hotspot Question

You have an on-premises Microsoft SQL Server instance named SQLSVR1 that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL database named SQLDB1.

You need to migrate the data stored in DB1 to SQLDB1 by using SQL Server replication. The solution must minimize the performance impact on DB1.

How should you configure the replication? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**



Configure SQLDB 1:

- As a push subscriber only
- As a push subscriber and a distributor
- To use a pull subscription only

Type of replication:

- Snapshot
- Transactional with peer-to-peer replication
- Transactional with updatable subscriptions

**Answer:**

**Answer Area**

Configure SQLDB 1:

- As a push subscriber only
- As a push subscriber and a distributor
- To use a pull subscription only

Type of replication:

- Snapshot
- Transactional with peer-to-peer replication
- Transactional with updatable subscriptions

Explanation:

Box 1: As a push subscriber and a distributor

To migrate a SQL Server database to Azure SQL Database using replication, you'll need to set up transactional replication with the on-premises SQL Server as the publisher and distributor and the Azure SQL Database as the subscriber. This process involves configuring the replication topology, including the publication, subscription, and the distributor. You'll also need to ensure proper authentication and firewall rules are set up.

Box 2: Transactional with updatable subscriptions

Snapshot and one-way transactional replication are supported. Peer-to-peer transactional replication and merge replication aren't supported.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/replication-to-sql-database>

**NEW QUESTION: 131**

Hotspot Question

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1. You need to ensure that app1 can access db1. The solution must use the principle of least privilege.


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

NOTE: Each correct selection is worth one point.

**Answer Area**

CREATE [app1]

CREDENTIAL	FOR LOGIN app1
LOGIN	FROM EXTERNAL PROVIDER
USER	FROM LOGIN app1
	WITHOUT LOGIN



**Answer:**

**Answer Area**

CREATE [app1]

CREDENTIAL	FOR LOGIN app1
LOGIN	FROM EXTERNAL PROVIDER
USER	FROM LOGIN app1
	WITHOUT LOGIN



Explanation:

"Create the service principal user in Azure SQL Database - Create the user AppSP in the SQL Database using the following T-SQL command:

```
CREATE USER [AppSP] FROM EXTERNAL PROVIDER
```

```
GO
```

```
"
```

Logical server is a default container for Azure SQL Database.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql>

**NEW QUESTION: 132**

### Hotspot Question

You have an Azure subscription.

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

- Support at least 36 cores.
- Minimize costs.

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

NOTE: Each correct selection is worth one point.

#### Answer Area

```
$SqlDatabaseParams = {  
    ResourceGroupName: "RG1",  
    ServerName: "Server01",  
    DatabaseName: "DB1",  
    Edition: "GeneralPurpose",  
    Vcore: 2,  
    ComputeGeneration:  
    ComputeModel:  
}  
  
New-AzSqlDatabase @SqlDatabaseParams
```

Microsoft

ComputeGeneration: 

"4"
"5"
"Gen4"
"Gen5"

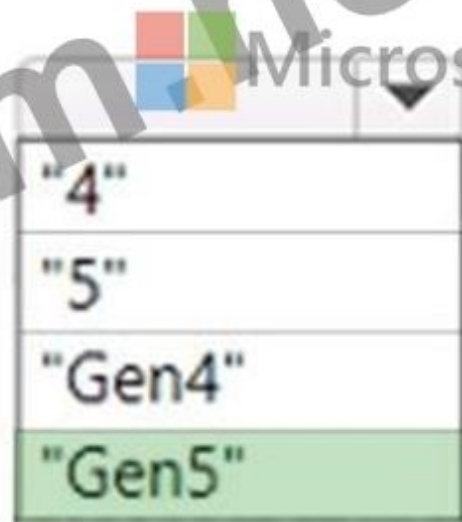
ComputeModel: 

Free
Provisioned
Serverless
Standard

**Answer:**

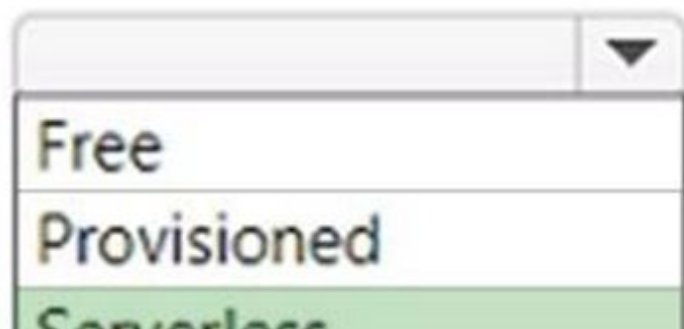
## Answer Area

```
$SqlDatabaseParams = {  
    ResourceGroupName: "RG1",  
    ServerName: "Server01",  
    DatabaseName: "DB1",  
    Edition: "GeneralPurpose",  
    Vcore: 2,  
    ComputeGeneration:
```



A dropdown menu for the 'ComputeGeneration' property. The menu is open, showing four options: "4", "5", "Gen4", and "Gen5". The "Gen5" option is currently selected and highlighted with a green background. A Microsoft logo is visible in the background of the dropdown.

```
    ComputeModel:
```



A dropdown menu for the 'ComputeModel' property. The menu is open, showing three options: "Free", "Provisioned", and "Serverless". The "Serverless" option is currently selected and highlighted with a green background.



Explanation:

Box 1: "Gen5"

Box 2: Serverless

Purchasing model and service tier

The following table describes serverless support based on purchasing model, service tiers, and hardware:

Category	Supported	Not supported
Purchasing model	vCore	DTU
Service tier	General Purpose Hyperscale	Business Critical
Hardware	Standard-series (Gen5)	All other hardware

The vCore purchasing model compute size is 2 to 128 vCores.

Reference:

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

### NEW QUESTION: 133

You have an Azure SQL database.

You need to identify a long-running query. Here you need to identify which operation in the question is causing a performance issue.

Which of the following can be used to display the query execution plan in Microsoft SQL Server Management Studio?

- A. Live Query Statistics
- B. An estimated execution plan
- C. An actual execution plan
- D. Client Statistics

**Answer: (SHOW ANSWER)**

With Live Query statistics, you get real-time information and insights on the queries as they are running. This allows you to debugging query performance issues, such as long-running queries, and queries that run indefinitely and never finish.

Option B is incorrect since this will only give you what could be the plan that is going to be used  
Option C is incorrect since this will not give live statistics that can show which query is causing an issue  
Option D is incorrect since this will not give the server stats that will show the performance of the queries

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure?view=sql-server-ver15>

### **NEW QUESTION: 134**

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

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

What should you do?

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

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

Automatic tuning is a fully managed intelligent performance service that uses built-in intelligence to continuously monitor queries executed on a database, and it automatically improves their performance.

Automatic tuning for Azure SQL Database uses the CREATE INDEX, DROP INDEX, and FORCE LAST GOOD PLAN database advisor recommendations to optimize your database performance.

Reference:

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

### **NEW QUESTION: 135**

You have an Azure subscription.

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

Which virtual machine series should you use?

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

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

<https://learn.microsoft.com/en-us/azure/virtual-machines/how-to-enable-write-accelerator>

### **NEW QUESTION: 136**

You have an on-premises server that has Microsoft SQL Server 2022 installed. The server contains multiple databases.

You have an Azure subscription that contains multiple Azure SQL databases.

You configure daily backups for the databases.

You need to ensure that all the backups are stored in a central location.

What should you use to store the backups?

- A. an Azure Files share
- B. Azure Table storage
- C. a Recovery Services Vault
- D. Azure Blob Storage

**Answer: (SHOW ANSWER)**

SQL Server backup and restore with Azure Blob Storage

To store Microsoft SQL Server 2022 database backups centrally, you can utilize features like Azure Blob Storage, S3-compatible object storage, or shared network drives. Consider using SQL Server Management Studio (SSMS) or Maintenance Plans to automate backups to these central locations.

Azure Blob Storage:

SQL Server 2022 integrates with Azure Blob Storage for backups, allowing you to store them in the cloud.

You can use the BACKUP TO URL syntax with the s3:// or https:// protocol to specify the Azure Blob Storage location.

You'll need to create an Azure storage account and a container, and then create a SQL Server credential to authenticate with the storage.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-and-restore-with-microsoft-azure-blob-storage-service>

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

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

<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

**Special Discount Code: freecram**)

### **NEW QUESTION: 137**

You have an Azure subscription that contains a Standard SKU Azure SQL database named DB1. You need to enable point-in-time recovery (PITR) backups for DB1. The solution must maximize how long the backups are retained.

Which retention period should you configure?

- A. seven days
- B. 35 days
- C. 180 days

D. 365 days

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

Backup storage (point-in-time restore)

By default, seven days of database backups are stored in RA-GRS Standard blob storage. The retention period can be increased to up to 35 days.

The Azure SQL Database service protects all databases with an automated backup system.

These backups are retained for 7 days for Basic, 35 days for Standard and 35 days for Premium.

Reference:

<https://azure.microsoft.com/en-us/pricing/details/azure-sql-database/single>

<https://azure.microsoft.com/en-us/blog/azure-sql-database-point-in-time-restore/>

### **NEW QUESTION: 138**

You have an Azure subscription that contains the following resources:

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

You need to implement a centralized monitoring solution for all the Azure SQL resources. The solution must minimize administrative effort.

What should you include in the solution?

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

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

### **NEW QUESTION: 139**

Hotspot Question

You have an Azure subscription.

You need to deploy a logical SQL server by using an Azure Resource Manager (ARM) template.

The solution must ensure that the server will allow inbound connectivity from any Azure resource.

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",  
    ...  
  },  
  {  
    "type": "Microsoft.Sql/servers/firewallRules",  
    "apiVersion": "2021-02-01-preview",  
    "name": "[concat(parameters('servers_a400m10102_name'), 'AllowAllWindowsAzureIps')]",  
    "dependsOn": [  
      "[resourceId('Microsoft.Sql/servers', parameters('servers_a400m10102_name'))]"  
    ],  
    "properties": {  
      "startIpAddress":   
      "endIpAddress":   
    }  
  }  
]
```

## Answer: Answer Area

```
"resources": [  
  {  
    "type": "Microsoft.Sql/servers",  
    ...  
  },  
  {  
    "type": "Microsoft.Sql/servers/firewallRules",  
    "apiVersion": "2021-02-01-preview",  
    "name": "[concat(parameters('servers_a400m10102_name'), 'AllowAllWindowsAzureIps')]",  
    "dependsOn": [  
      "[resourceId('Microsoft.Sql/servers', parameters('servers_a400m10102_name'))]"  
    ],  
    "properties": {  
      "startIpAddress":   
      "endIpAddress":   
    }  
  }  
]
```

Explanation:

[https://learn.microsoft.com/en-](https://learn.microsoft.com/en-us/azure/templates/microsoft.sql/servers/firewallrules?pivots=deployment-language-bicep)

[us/azure/templates/microsoft.sql/servers/firewallrules?pivots=deployment-language-bicep](https://learn.microsoft.com/en-us/azure/templates/microsoft.sql/servers/firewallrules?pivots=deployment-language-bicep) The

start IP address of the firewall rule. Must be IPv4 format. Use value '0.0.0.0' for all Azure- internal IP addresses.

#### NEW QUESTION: 140

You have a Microsoft SQL Server 2022 instance.

You need to troubleshoot failing SQL Agent Jobs.

Which command should you run?

- A. EXEC sp\_trace\_setfilter
- B. EXEC sp\_help\_jobserver
- C. EXEC sp\_help\_jobsteplog
- D. EXEC sp\_read\_error\_log;

**Answer: (SHOW ANSWER)**

sp\_readerrorlog (Transact-SQL) allows you to read the contents of the SQL Server or SQL Server Agent error log file and filter on keywords.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-readerrorlog-transact-sql?>

#### NEW QUESTION: 141

You have an on-premises multi-tier application named App1 that includes a web tier, an application tier, and a Microsoft SQL Server tier. All the tiers run on Hyper-V virtual machines.

Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure.

You need to recommend a solution to fail over the database tier of App1 to Azure. The solution must provide the ability to test failover to Azure without affecting the current environment.

What should you include in the recommendation?

- A. Azure Backup
- B. Azure Information Protection
- C. Windows Server Failover Cluster
- D. Azure Site Recovery

**Answer: (SHOW ANSWER)**

You can use Azure Site Recovery to test failover.

You run a test failover to validate your replication and disaster recovery strategy, without any data loss or downtime. A test failover doesn't impact ongoing replication, or your production environment. You can run a test failover on a specific virtual machine (VM), or on a recovery plan containing multiple VMs.

Reference:

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure>

#### NEW QUESTION: 142

You have an Azure SQL managed instance named MI1.

You need to implement automatic tuning for the databases of MI1.

What should you do?

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

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

For Azure SQL Managed Instance, the supported option FORCE\_LAST\_GOOD\_PLAN can only be configured through T-SQL. The Azure portal based configuration and automatic index tuning options described in this article do not apply to Azure SQL Managed Instance.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable?view=azuresql>

### NEW QUESTION: 143

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1. You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1. and a virtual network named VNET1. SQLMI1 resides on VNET1. The on-premises network connects to VNET1 by using an ExpressRoute connection.

You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service.

You need to configure VNET1 to support the migration.

What should you do?

- A. Configure service endpoints.
- B. Configure virtual network peering.
- C. Deploy an Azure firewall.
- D. Configure network security groups (NSGs).

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

During virtual network setup, if you use ExpressRoute with network peering to Microsoft, add the following service endpoints to the subnet in which the service will be provisioned:

- \* Target database endpoint (for example, SQL endpoint, Cosmos DB endpoint, and so on)
- \* Storage endpoint
- \* Service bus endpoint

This configuration is necessary because Azure Database Migration Service lacks internet connectivity.

Reference:

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

### NEW QUESTION: 144

You have an Azure Data Factory that contains 10 pipelines.

You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.

What should you add to each pipeline?

- A. an annotation

- B. a resource tag
- C. a run group ID
- D. a user property
- E. a correlation ID

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

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag.

You can define tags so you can see their performance or find errors faster.

Reference:

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

### NEW QUESTION: 145

Hotspot Question

You have an Azure subscription that contains an Azure SQL managed instance named SQL1.

You use the Microsoft Power BI service.

You need to configure connectivity from the Power BI service to SQL1. The solution must ensure that only the Power BI service can initiate connections to SQL1.

Which type of endpoint should you use for SQL1, and what should you use to control the connections to SQL1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Endpoint type: 

	▼
Private	
Public	
Service	

Control connections by using: 

	▼
A database-level firewall	
A network security group (NSG)	
A server-level firewall	
Virtual network firewall rules	

**Answer:**

Endpoint type:

	▼
Private	
Public	
Service	

Control connections by using:

	▼
A database-level firewall	
A network security group (NSG)	
A server-level firewall	
Virtual network firewall rules	

Explanation:

To configure connectivity from the Power BI service to an Azure SQL managed instance, you should use a Public endpoint for SQL1 1. You can control the connections to SQL1 by using a Network Security Group (NSG) 1. Therefore, the correct answer is Endpoint type: Public and Control connections by using: A network security group (NSG).

### NEW QUESTION: 146

Case Study 1 - Litware, Inc

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.

### Drag and Drop Question

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.

The screenshot shows a drag-and-drop question interface. On the left, under the heading "Actions", there are six items in a list:

- 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

On the right, under the heading "Answer Area", there are two empty slots. Each slot has a circular icon with an arrow pointing up or down, indicating that the order of the selected actions matters. A watermark "freedam.net Microsoft" is visible across the interface.

**Answer:**

## Actions

Create a user in the master database

Modify the Azure SQL server administrator account

## 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

Enable the contained database authentication option

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.

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

Reference:

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

## NEW QUESTION: 147

Hotspot Question

You have a Microsoft SQL Server 2017 server that hosts five databases.

You Plan to migrate the databases to Azure.

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


- Automatically scales compute based on the workload demand
- Provides per-second billing

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

**Answer Area**

Azure service:


Service tier:



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

Azure service:

Service tier:



**NEW QUESTION: 148**

You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:

- Support failovers that do not require client applications to change their connection strings.
- Replicate the database to a secondary Azure region.
- Support failover to the secondary region.

What should you include in the recommendation?

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

**Answer: (SHOW ANSWER)**

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).

Incorrect Answers:

C: Availability Zones are unique physical locations within a region. Each zone is made up of one

or more datacenters equipped with independent power, cooling, and networking.

Reference:

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

**NEW QUESTION: 149**

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)**

SQL Server Management Studio provides the ability to view the live execution plan of an active query. This live query plan provides real-time insights into the query execution process as the controls flow from one query plan operator to another. The live query plan displays the overall query progress and operator-level run-time execution statistics such as the number of rows produced, elapsed time, operator progress, etc. Because this data is available in real time without needing to wait for the query to complete, these execution statistics are extremely useful for debugging query performance issues.

Reference:

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

**NEW QUESTION: 150**

Hotspot Question

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.

**Answer Area**

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.



▼

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"

▼

a header
a data
an empty

**Answer:**

**Answer Area**

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.

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"

a header
a data
an empty

Explanation:

Box 1: CSV files that have file named beginning with "tripdata\_2020"

Box 2: a header

FIRSTROW = 'first\_row'

Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Reference:

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

### NEW QUESTION: 151

Hotspot Question

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 Area

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE

Value:

1
60
CUSTOM
ON

Answer:

## Answer Area

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE

Value:

1
60
CUSTOM
ON



Explanation:

#### Box 1: INTERVAL\_LENGTH\_MINUTES

INTERVAL\_LENGTH\_MINUTES defines size of time window during which collected runtime statistics for query plans are aggregated and persisted. Every active query plan has at most one row for a period of time defined with this configuration.

Default: 60

#### Box 2: 1

Statistics Collection Interval (INTERVAL\_LENGTH\_MINUTES): Defines the level of granularity for the collected runtime statistic, expressed in minutes. The default is 60 minutes. Consider using a lower value if you require finer granularity or less time to detect and mitigate issues.

Use SQL Server Management Studio or Transact-SQL to set a different value for Statistics Collection Interval:

```
ALTER DATABASE [QueryStoreDB]
SET QUERY_STORE (INTERVAL_LENGTH_MINUTES = 60);
```

Reference:

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

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdisscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

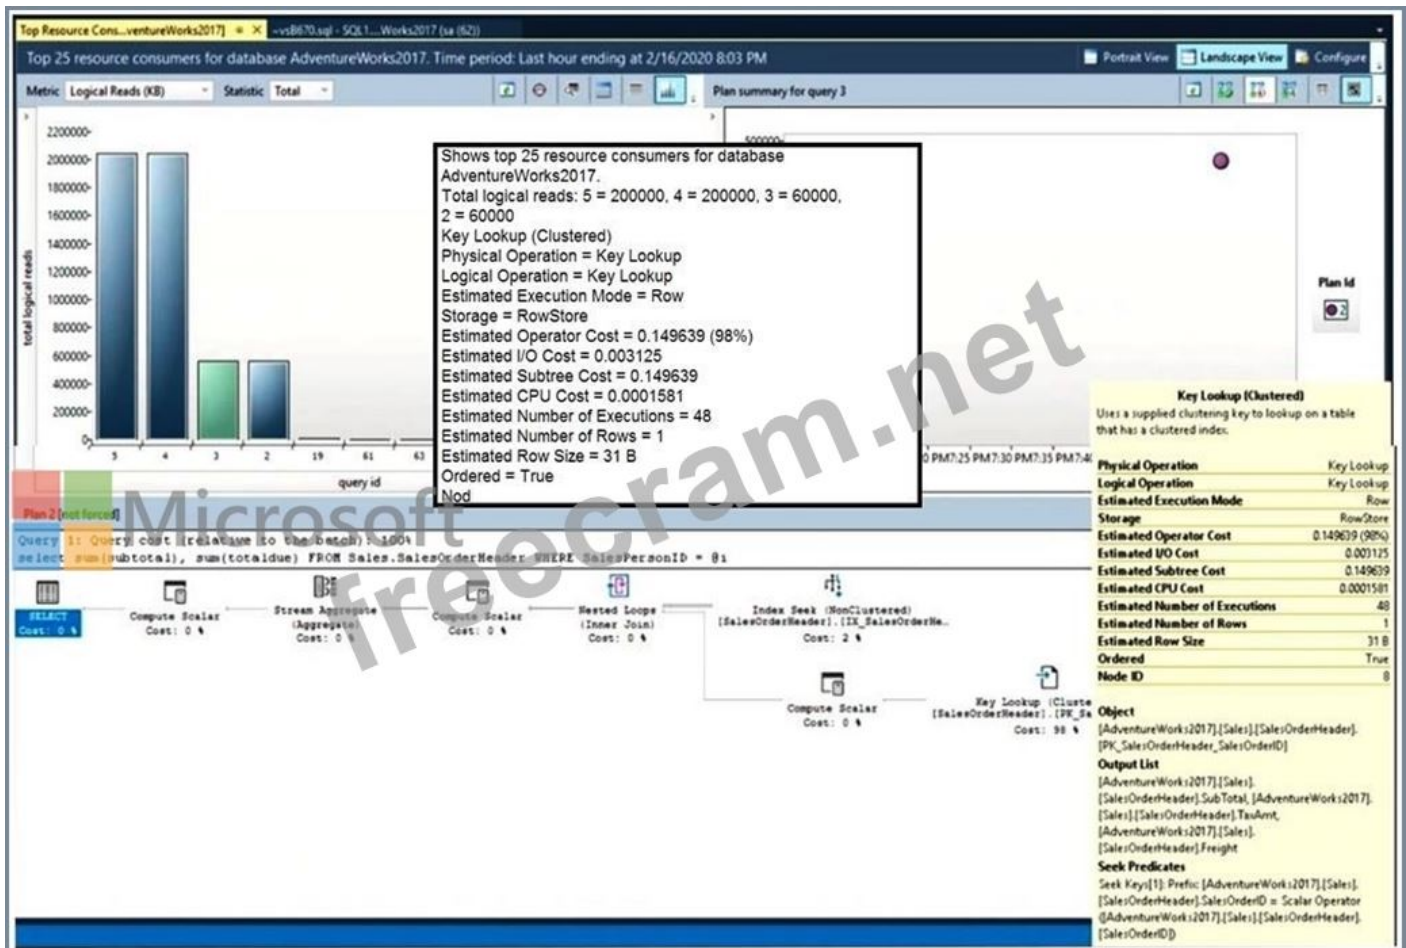
**Special Discount Code: **freecram****)

#### **NEW QUESTION: 152**

Hotspot Question

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.

**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		Microsoft	
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 checked="" 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:

Box 1: No

There is only one query plan available. Force has no effect.

Box 2: No

Adding an index will not increase IO usage.

Box 3: Yes

The performance would improve.

Reference:

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

### NEW QUESTION: 153

You have an Azure subscription.

You plan to deploy an Azure SQL managed instance named Server1.

You need to recommend a scaling solution for Server1. The solution must ensure that Server1 can dynamically scale reads and writes.

What should you include in the recommendation?

- A. a vCore based service tier
- B. elastic pools
- C. A Hyperscale based service tier
- D. a DTU based service tier

**Answer: (SHOW ANSWER)**

Dynamically scale database resources with minimal downtime - Azure SQL Database & Azure SQL Managed Instance Azure SQL Managed Instance allows you to scale as well:

SQL Managed Instance uses vCores mode and enables you to define maximum CPU cores and maximum of storage allocated to your instance. All databases within the managed instance will share the resources allocated to the instance.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/scale-resources>

### NEW QUESTION: 154

## SIMULATION

You need to add an Azure AD user named user2-12345678@examusers.com to db1. User2-12345678 must be able to read data from all the tables in db1 without being able to modify the data.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

### **Answer:**

Stage 1: Create the user

Create a user with SSMS (SQL Server Management Studio)

Step 1: In Object Explorer, expand the Databases folder.

Step 2: Expand the database in which to create the new database user. [Here db1] Step 3: Right-click the Security folder, point to New, and select User....

Step 4: In the Database User - New dialog box, on the General page, select one of the following user types from the User type list:

SQL user with login

SQL user with password (when contained database is enabled)

SQL user without login

User mapped to a certificate

User mapped to an asymmetric key

\*-> Windows user

Step 4a: Select Windows User

Step 5: When you select an option, the remaining options in the dialog may change. Some options only apply to specific types of database users. Some options can be left blank and will use a default value.

User name

Enter a name for the new user. If you have chosen Windows user from the User type list, you can also select the ellipsis (...) to open the Select User or Group dialog box. [Specify User2-12345678]

Step 6: Select OK.

Stage 2: Assign permissions

Additional Options

The Database User - New dialog box also offers options on four other pages: Owned Schemas, Membership, Securables, and Extended Properties.

Step 7: Select Securables

The Securables page lists all possible securables and the permissions on those securables that can be granted to the login.

Step 8: Select tables

Step 9: Grant Read Permission

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/security/authentication-access/create-a-database-user>

**NEW QUESTION: 155**

You have an Azure subscription that contains an Azure SQL database named DB1. You need to recommend a high availability, disaster recovery (HA/DR) solution for DB1. The solution must meet the following requirements:

- Ensure that a recovery time objective (RTO) of less than five minutes can be met.
- Ensure that a recovery point objective (RPO) of less than 30 seconds can be met.
- Ensure availability in the event of an Azure region outage.

What should you include in the recommendation?

- A.** zone redundancy
- B.** failover groups
- C.** geo-restore
- D.** Azure Site Recovery

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

To protect an Azure SQL Database from a regional outage, you should implement disaster recovery solutions like active geo-replication or failover groups to create a secondary, readable replica in a different Azure region.

Azure failover groups for Azure SQL Database offer a Recovery Time Objective (RTO) of typically less than 60 seconds and a Recovery Point Objective (RPO) of 5 seconds. This means that, in the event of a failover, your database should be back online within a minute, and you might experience a maximum of 5 seconds of data loss.

Note:

Active Geo-Replication and Failover Groups:

Active Geo-Replication:

This feature allows you to create readable secondary databases in different Azure regions. It provides a mechanism for failover to the secondary database in case of a regional outage, minimizing downtime.

Failover Groups:

Failover groups simplify the management of geo-replicated databases, enabling automatic failover of multiple databases simultaneously. They also provide read-write and read-only listener endpoints that remain unchanged after failover, simplifying application updates.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/failover-group-sql-db?view=azuresql>

**NEW QUESTION: 156**

You have an Azure subscription.

You create a logical SQL server that hosts four databases. Each database will be used by a separate customer.

You need to ensure that each customer can access only its own database. The solution must minimize administrative effort. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Create a database-level firewall rule.
- B. Deny public access.
- C. Create a private endpoint.
- D. Create a server-level firewall rule.
- E. Create a network security group (NSG).

Answer: ([SHOW ANSWER](#))

### NEW QUESTION: 157


Hotspot Question

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

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

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area** 

Backup type:  ▼

- Full and log backups only
- Full backup only
- Log backup only

Backup option:  ▼

- WITH CHECKSUM
- WITH NOINIT
- WITH UNLOAD

Answer:


## Answer Area

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD



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: 158

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2. You need to migrate the databases hosted on SQL1 to Azure. The solution must meet the following requirements:

- The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.
- Administrative effort must be minimized.

What should you use to host the databases?

- A. a single Azure SQL database
- B. an Azure SQL Database elastic pool
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed Instance

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

Linked servers are available in SQL Server Database Engine and Azure SQL Managed Instance. They are not enabled in Azure SQL Database singleton and elastic pools.

<https://learn.microsoft.com/en-us/sql/relational-databases/linked-servers/linked-servers-database-engine?view=sql-server-2017>

**NEW QUESTION: 159**

Your environment contains a Microsoft Azure virtual machine and an Azure SQL database. The virtual machine has Microsoft SQL Server installed and contains several applications.

You need to recommend a solution to migrate one of the applications to the Azure SQL database.

The solution must minimize downtime.

Which solution should you recommend?

- A. Implement log shipping.
- B. Implement transactional replication.
- C. Restore a BACPAC of the database.
- D. Restore a backup of the database.

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

**NEW QUESTION: 160**

Hotspot Question

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 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:

## 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	

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: 161

You deploy an instance of SQL Server on Azure Virtual Machines: named SQL1 that hosts multiple databases.

You configure the full recovery model for all the databases.

You perform a full backup of the master database on SQL1.

You need to perform an additional backup of the master database on SQL1. The solution must minimize how long it takes to perform the backup.

Which type of backup should you perform?

- A. log
- B. full
- C. differential
- D. tail-log

**Answer: (SHOW ANSWER)**

Backups, only a full database backup can be performed on the master database.

<https://learn.microsoft.com/en-us/sql/relational-databases/databases/master-database?view=sql-server-ver16>

**NEW QUESTION: 162**

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: (SHOW ANSWER)**

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>

## **NEW QUESTION: 163**

Case Study 1 - Litware, Inc

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.

#### Hotspot Question

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 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



Explanation:

Box 1: Premium 4-VCORE

Scenario:

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

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 virtual network that has service endpoints

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

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.

Note: During virtual network setup, if you use ExpressRoute with network peering to Microsoft, add the following service endpoints to the subnet in which the service will be provisioned:

Target database endpoint (for example, SQL endpoint, Cosmos DB endpoint, and so on) Storage endpoint Service bus endpoint This configuration is necessary because Azure Database Migration Service lacks internet connectivity.

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: 164

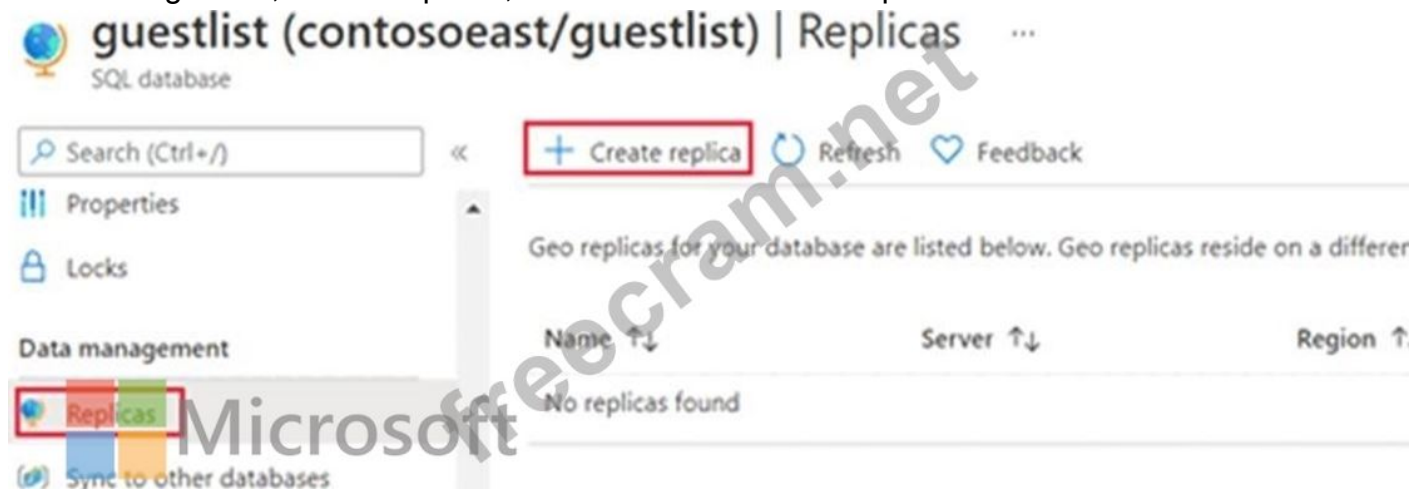
### SIMULATION

You need to implement a disaster recovery solution by using active geo-replication for an Azure SQL database named db1. The replica must be in the East US Azure region on a server named sql23456789-dr.database.windows.net.

You may need to use SQL Server Management Studio and the Azure portal.

### Answer:

Step 1: Navigate to your database: In the Azure portal, browse to the Azure SQL Database you want to replicate. [Select db1] Step 2: On the SQL Database page, select your database, scroll to Data management, select Replicas, and then select Create replica.



The screenshot shows the Azure portal interface for an Azure SQL database. The breadcrumb navigation path is 'guestlist (contosoeast/guestlist) | Replicas'. The page title is 'SQL database'. On the left sidebar, the 'Data management' section is expanded, and the 'Replicas' option is selected. The main content area shows a '+ Create replica' button highlighted with a red box. Below the button, there is a table header with columns 'Name', 'Server', and 'Region', and the text 'No replicas found'.

Step 3: Select your geo-secondary database Subscription and Resource group.

**Geo-secondary database details**

**Project details**

Select the subscription of geo secondary database to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription \* ⓘ

Resource group \* ⓘ  [Create new](#)

Step 4: Select or create the server for the secondary database [Select named sql23456789-dr.database.windows.net] Question: The replica must be in the East US Azure region on a server named sql23456789- dr.database.windows.net.

Step 5: When secondary database is created [or selected], the deployment process begins.

Delete Cancel Redeploy Refresh

**Deployment is in progress**

Deployment name: Microsoft.SQLDatabase.newDatabaseCopyExisti... Start time: 8/20/2021, 3:16:00 PM  
 Subscription: Contoso Correlation ID:  
 Resource group: ContosoHotel

Deployment details (Download)

Resource	Type	Status	Operation details
contosowest/guestlist	Microsoft.Sql/servers/databases	Accepted	<a href="#">Operation details</a>

Step 6: When the deployment is complete, the secondary database displays its status.

**Your deployment is complete**

Deployment name: Microsoft.SQLDatabase.newDatabaseCopyExisti... Start time: 8/20/2021, 3:16:00 PM  
 Subscription: Contoso Correlation ID:  
 Resource group: ContosoHotel

Deployment details (Download)

Resource	Type	Status	Operation details
contosowest/guestlist	Microsoft.Sql/servers/databases	Created	<a href="#">Operation details</a>

(Step 7: Return to the primary database page, and then select Replicas. Your secondary database is listed under Geo replicas.) Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-configure-portal>

**NEW QUESTION: 165**

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

## Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

## Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

## Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

## Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1. App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event\_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

## Requirements

### Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

### Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.

- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

#### Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

#### Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

#### Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

You need to recommend a solution that will enable remote developers to access DB1 and DB2.

The solution must support the planned changes and meet the security requirements.

What should you include in the recommendation?

- A. a public endpoint via a database-level firewall rule
- B. a Point-to-Site (P2S) VPN
- C. a public endpoint via a server-level firewall rule
- D. a private endpoint

**Answer: (SHOW ANSWER)**

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/connect-application-instance?view=azuresql#connect-from-inside-a-different-vnet>

#### **NEW QUESTION: 166**

Hotspot Question

You configure a long-term retention policy for an Azure SQL database as shown in the exhibit.

## 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?

Monthly LTR Backups ⓘ

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

Yearly LTR Backups ⓘ

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

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

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	<input type="text"/>
	6 weeks
	12 months
	10 years
The backup saved to long-term retention on January 11, 2020 will be retained for	<input type="text"/>
	6 weeks
	12 months
	10 years

**Answer:**

The backup saved to long-term retention on January 4, 2020, will be retained for	<input type="text"/>
	6 weeks
	12 months
	10 years
The backup saved to long-term retention on January 11, 2020 will be retained for	<input type="text"/>
	6 weeks
	12 months
	10 years

Explanation:

Box 1: 12 months

The January 4, 2020 backup is the first backup of the month, and it will keep for 12 months.

Box 2: 10 years.

The January 11, 2020 backup is the Week 2 backup. This backup will be the yearly backup, and it will be retained for 10 years.

Note: Long-term backup retention (LTR) leverages the full database backups that are automatically created to enable point in time restore (PITR). If an LTR policy is configured, these backups are copied to different blobs for long-term storage. The copy is a background job that has no performance impact on the database workload. The LTR policy for each database in SQL Database can also specify how frequently the LTR backups are created.

To enable LTR, you can define a policy using a combination of four parameters: weekly backup retention (W), monthly backup retention (M), yearly backup retention (Y), and week of year (WeekOfYear). If you specify W, one backup every week will be copied to the long-term storage. If you specify M, the first backup of each month will be copied to the long-term storage. If you specify Y, one backup during the week specified by WeekOfYear will be copied to the long-term storage. If the specified WeekOfYear is in the past when the policy is configured, the first LTR backup will be created in the following year. Each backup will be kept in the long-term storage according to the policy parameters that are configured when the LTR backup is created.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/long-term-retention-overview?view=azuresql>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:  
<https://www.examdisscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

### NEW QUESTION: 167

You have an Azure subscription that contains three instances of SQL Server on Azure Virtual Machines.

You plan to implement a disaster recovery solution.

You need to be able to perform disaster recovery drills regularly. The solution must meet the following requirements:

- Minimize administrative effort for the recovery drills.
- Isolate the recovery environment from the production environment

What should you use?

- A. Azure Site Recovery
- B. Recovery Services vaults
- C. Azure Backup
- D. native Microsoft SQL Server backup

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

### NEW QUESTION: 168

Drag and Drop Question

You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.

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

Actions	Answer Area
Duplicate Azure Policy definitions	
Run Azure Policy remediation tasks	
Create an Azure Blueprints assignment	
Create an Azure Policy initiative	
Create an Azure Policy initiative assignment	

**Answer:**

Actions	Answer Area
Duplicate Azure Policy definitions	Create an Azure Policy initiative
	Create an Azure Policy initiative assignment
Create an Azure Blueprints assignment	Run Azure Policy remediation tasks

**Explanation:**

**Step 1: Create an Azure Policy Initiative**

The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.

With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.

**Step 2: Create an Azure Policy Initiative assignment**

Assign the initiative definition you created in the previous step.

**Step 3: Run Azure Policy remediation tasks**

To apply the Policy Initiative to the existing SQL databases.

**Reference:**

<https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage>

**NEW QUESTION: 169**

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: (SHOW ANSWER)**

You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the Get Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity>

### NEW QUESTION: 170

Your on-premises network contains a server that hosts a 60-TB database named DB1. The network has a 10-Mbps internet connection.

You need to migrate DB1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

A. Data Migration Assistant (DMA)

B. Azure Migrate

C. Azure Data Studio

D. Azure Data Box

**Answer: (SHOW ANSWER)**

Microsoft Azure Data Box is a hardware appliance designed to allow customers to import or export large amounts of data - more than 40TB - into and out of Azure offline. It is especially helpful when there is zero or limited network connectivity. Microsoft ships customers a proprietary Data Box storage device with a rugged casing to protect and secure data during the transit. A customer may choose Data Box for a one-time or the occasional cloud migration or an initial bulk data transfer followed by periodic transfers.

### NEW QUESTION: 171

You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image. You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance.

What should you do first?

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

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

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/change-sql-server-edition?view=azuresql&tabs=azure-portal>

### NEW QUESTION: 172

Drag and Drop Question

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.

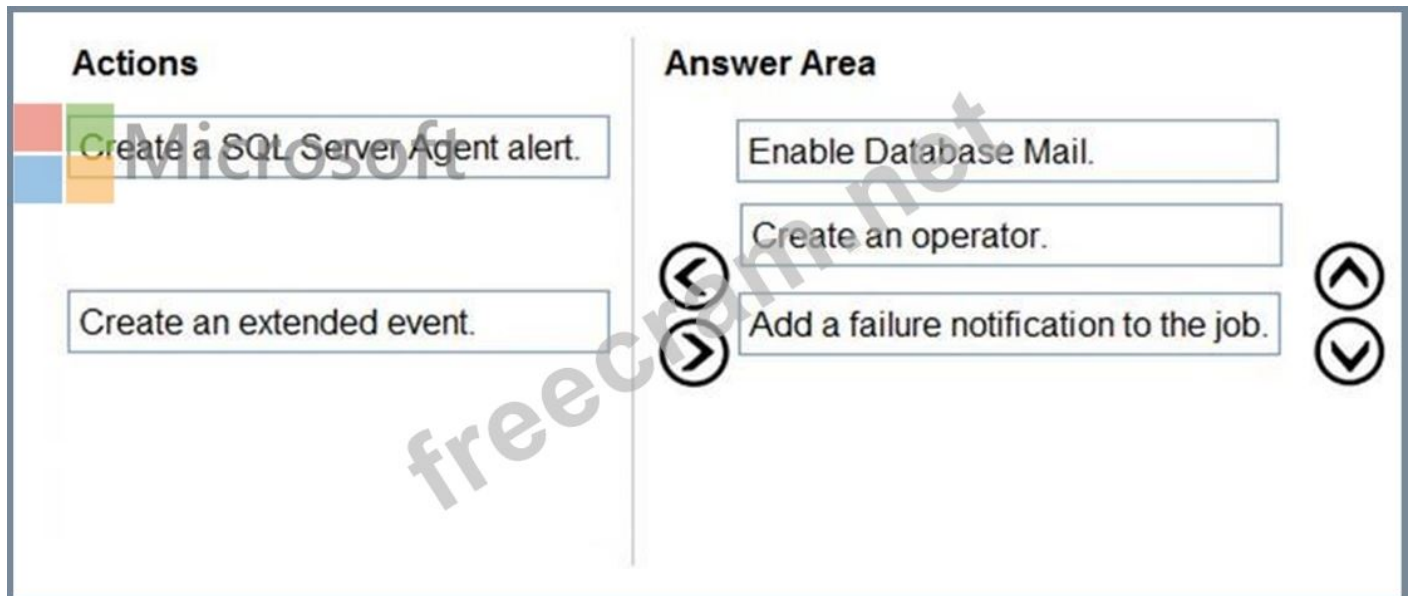
You need to be notified by email if the job fails.

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.

The screenshot shows a drag-and-drop interface with two main sections: 'Actions' and 'Answer Area'. The 'Actions' section contains five items in a list: 'Create a SQL Server Agent alert.', 'Create an operator.', 'Create an extended event.', 'Enable Database Mail.', and 'Add a failure notification to the job.'. The 'Answer Area' contains four arrow icons: a left arrow, a right arrow, an up arrow, and a down arrow, used for moving and ordering the items. A watermark 'freecram.net' is overlaid on the interface, and the Microsoft logo is visible at the bottom.

**Answer:**



Explanation:

Step 1: Enable Database Mail

If it isn't already enabled, first you would need to configure the Database Mail feature on SQL Managed Instance.

Box 2: Create an operator.

You can notify the operator that something happened with your SQL Agent jobs. An operator defines contact information for an individual responsible for the maintenance of one or more instances in SQL Managed Instance.

Box 3: Add a failure notification to the job,

You can then modify any SQL Agent job and assign operators that will be notified via email if the job completes, fails, or succeeds using SSMS or the following T-SQL script:

```
EXEC msdb.dbo.sp_update_job @job_name=N'Load data using SSIS',  
@notify_level_email=3, -- Options are: 1 on succeed, 2 on failure, 3 on complete  
@notify_email_operator_name=N'AzureSQLTeam';
```

Reference:

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

### NEW QUESTION: 173

You have SQL Server on an Azure virtual machine.

You need to add a 4-TB volume that meets the following requirements:

- Maximizes IOPs
- Uses premium solid state drives (SSDs)

What should you do?

- A. Attach two mirrored 4-TB SSDs.
- B. Attach a stripe set that contains four 1-TB SSDs.
- C. Attach a RAID-5 array that contains five 1-TB SSDs.
- D. Attach a single 4-TB SSD.

Answer: ([SHOW ANSWER](#))

For more throughput, you can add additional data disks and use disk striping.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/storage-configuration?tabs=windows2016>

### NEW QUESTION: 174

Hotspot Question

You have an Azure subscription.

You plan to provision a logical SQL server for Azure SQL Database by using a Bicep template.

You need to configure the resource section of the template to ensure that the deployment always maximizes performance for network connections to the server, regardless of whether the connection originates from within or outside of Azure.

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

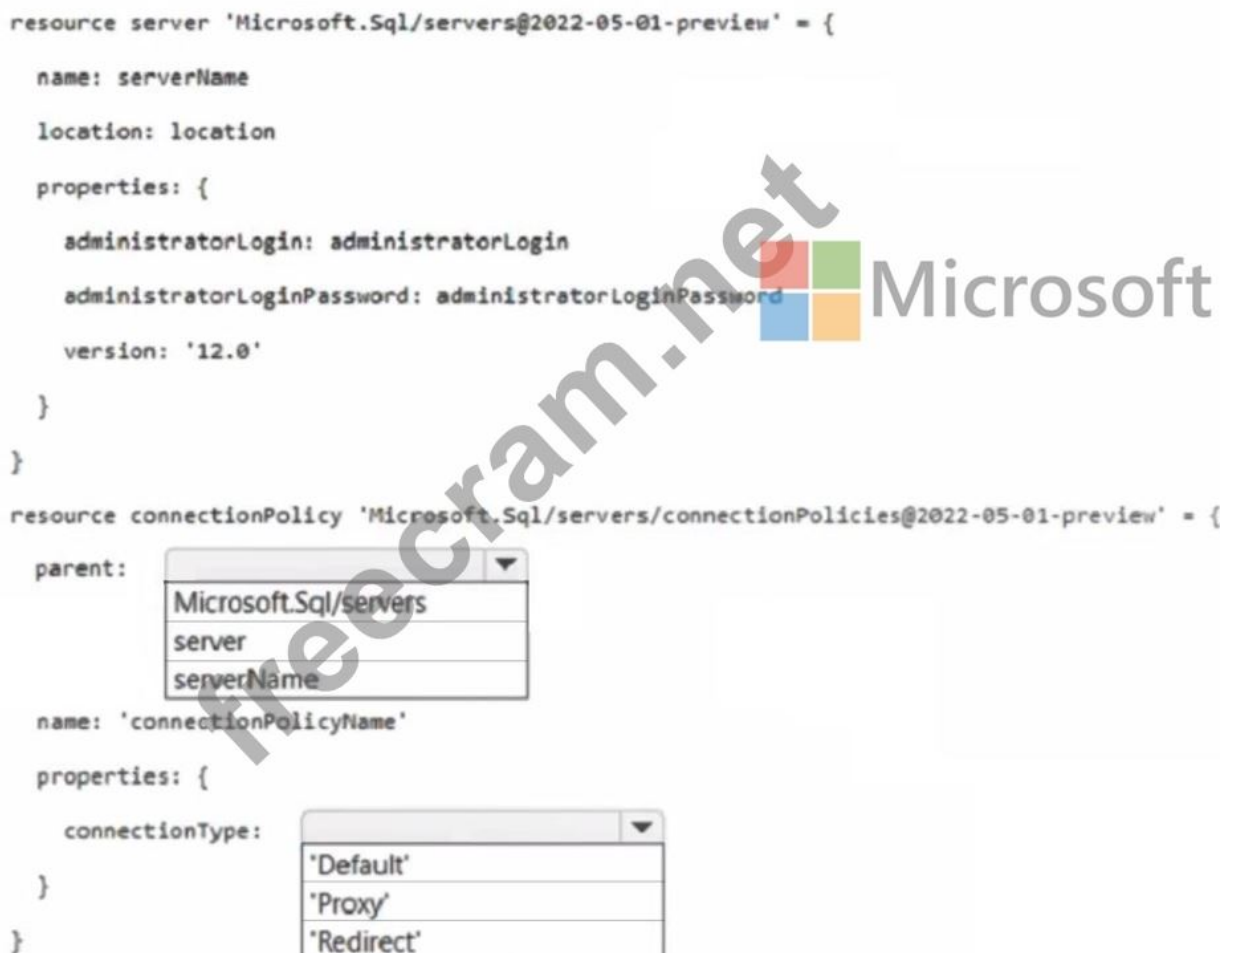
```
resource server 'Microsoft.Sql/servers@2022-05-01-preview' = {  
  name: serverName  
  location: location  
  properties: {  
    administratorLogin: administratorLogin  
    administratorLoginPassword: administratorLoginPassword  
    version: '12.0'  
  }  
}  
  
resource connectionPolicy 'Microsoft.Sql/servers/connectionPolicies@2022-05-01-preview' = {  
  parent: 

|                       |
|-----------------------|
| Microsoft.Sql/servers |
| server                |
| serverName            |

  
  name: 'connectionPolicyName'  
  properties: {  
    connectionType: 

|            |
|------------|
| 'Default'  |
| 'Proxy'    |
| 'Redirect' |

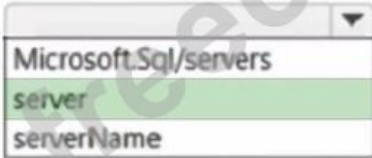
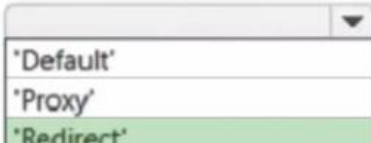
  
  }  
}
```



Answer:

```

resource server 'Microsoft.Sql/servers@2022-05-01-preview' {
  name: serverName
  location: location
  properties: {
    administratorLogin: administratorLogin
    administratorLoginPassword: administratorLoginPassword
    version: '12.0'
  }
}

resource connectionPolicy 'Microsoft.Sql/servers/connectionPolicies@2022-05-01-preview' =
  parent: 
  name: 'connectionPolicyName'
  properties: {
    connectionType: 
  }
}

```

## NEW QUESTION: 175

### SIMULATION

You need to increase the maximum storage size of the query store on db1 to 4,000 MB.

You may need to use SQL Server Management Studio and the Azure portal.

### Answer:

How to increase the maximum storage size of the query store for an azure database using Azure portal?

How to increase the maximum storage size of the query store for an azure database using SQL Server Management Studio?

To increase the size of Azure Query Store, you can use Transact-SQL within either SSMS (SQL Server Management Studio) or the Azure portal. Specifically, you'll need to modify the `MAX_STORAGE_SIZE_MB` setting for the database. No GUI options exist for direct modification within the Azure portal.

Step 1: Connect to your database:

Use SSMS or the Azure portal to connect to the database where you want to increase the Query Store size. [Connect to db1] Step 2: . Execute the ALTER DATABASE statement:

Run the following Transact-SQL query (replacing YourDB with your database name and 1024 with the desired storage size in MB):

```
ALTER DATABASE [YourDB] SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 1024);
```

\*-> Use the following query: ALTER DATABASE db1 SET QUERY\_STORE (MAX\_STORAGE\_SIZE\_MB = 4000);

Step 3: Apply changes: The Query Store will automatically adjust to the new size.

Reference:

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

### NEW QUESTION: 176

Hotspot Question

You have an Azure SQL database named DB1.

You need to identify how much unused space in megabytes was allocated to DB1.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

```
SELECT SUM(size/ [16.0] - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
AS DatabaseDataSpaceAllocatedUnusedInMB
FROM [sys.database_files]
GROUP BY type_desc
HAVING type_desc = 'ROWS';
```

Microsoft

Answer:

**Answer Area**

```
SELECT SUM(size/ [128.0] - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
AS DatabaseDataSpaceAllocatedUnusedInMB
FROM [sys.database_files]
GROUP BY type_desc
HAVING type_desc = 'ROWS';
```

Microsoft

Explanation:

Box 1: 128.0

-- Connect to database

-- Database data space allocated in MB and database data space allocated unused in MB

```
SELECT SUM(size/128.0) AS DatabaseDataSpaceAllocatedInMB, SUM(size/128.0 -  
CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0) AS
```

```
DatabaseDataSpaceAllocatedUnusedInMB FROM sys.database_files GROUP BY type_desc  
HAVING type_desc = 'ROWS' Box 2: sys.database_files Reference:
```

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/how-to-get-azure-sql-database-size/ba-p/369189>

### NEW QUESTION: 177

Drag and Drop Question

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



Microsoft



### Answer:

#### Actions

Export a built-in policy.

View the policy history.

Import a policy file.

#### Answer Area

Create a custom condition based on a built-in facet.

Create a custom policy based on a condition.

Run a policy evaluation.



Explanation:

First create a condition, then a custom policy based on the condition, finally run a policy evaluation.

Reference:

<https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-based-management/>

### NEW QUESTION: 178

Hotspot Question

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.

NOTE: Each correct selection is worth one point.

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 Microsoft

Statements	Yes	No
The log file will be truncated.	<input type="radio"/>	<input checked="" 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 checked="" type="radio"/>	<input type="radio"/>

Explanation:

Box 1: No

A copy-only log backup preserves the existing log archive point and, therefore, does not affect the sequencing of regular log backups.

Box 2: No

Box 3: Yes

If the database is online and you plan to perform a restore operation on the database, begin by backing up the tail of the log. To avoid an error for an online database, you must use the ... WITH NORECOVERY option of the BACKUP Transact-SQL statement.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/tail-log-backups-sql-server>

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/copy-only-backups-sql-server>

## NEW QUESTION: 179

You have an Azure Database for MySQL database. You have to identify which database queries are consuming the most resources.

Which of the following tool can be used for this requirement?

- 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: 180

You have 50 Azure SQL databases.

You need to automate Transact-SQL queries to run periodically against each database. The solution must minimize administrative effort.

What should you configure?

- A. elastic jobs
- B. SQL Server Agent jobs
- C. Power Automate
- D. Task Scheduler in Windows

**Answer: (SHOW ANSWER)**

The SQL Server Agent is a Microsoft Windows service that executes scheduled administrative tasks (called jobs) in SQL Server and Azure SQL Managed Instance.

Reference:

<https://learn.microsoft.com/en-us/ssms/agent/sql-server-agent>

### NEW QUESTION: 181

You have an Azure subscription.

You plan to deploy an instance of SQL Server on Azure Virtual Machines by using an Azure Marketplace image.

You need to register the SQL Server IaaS Agent extension (SqlIaaSExtension). The solution must meet the following requirements:

- Install critical updates for SQL Server automatically.
- Minimize performance impact on the virtual machine.

Which management mode should you select?

- A. full
- B. lightweight
- C. NoAgent

**Answer: (SHOW ANSWER)**

Full mode installs the SQL IaaS Agent to the VM to deliver full functionality. Use it for managing a SQL Server VM with a single instance. Full mode installs two Windows services that have a minimal impact to memory and CPU.

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

agent-extension-automate-management?view=azuresql&tabs=azure-powershell#management-modes

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam! ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examd Discuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

**Special Discount Code: **freecram****)

### **NEW QUESTION: 182**

You have the following resources:

- 15 SQL Server on Azure Virtual Machines instances
- 20 Azure SQL databases

You need to recommend a solution to centrally monitor the resources for security vulnerabilities.

What should you include in the recommendation?

- A. database audits
- B. Microsoft Defender
- C. SQL insights
- D. Azure SQL Auditing

**Answer: (SHOW ANSWER)**

<https://learn.microsoft.com/en-us/azure/defender-for-cloud/defender-for-sql-introduction>

### **NEW QUESTION: 183**

You have an Azure SQL database named SQL1.

You need to implement a disaster recovery solution for SQL1. The solution must minimize the following:

- The recovery point objective (RPO)
- The recovery time objective (RTO)
- Administrative effort

What should you include in the solution?

- A. Azure Site Recovery
- B. active geo-replication
- C. availability groups
- D. auto-failover groups

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

<https://learn.microsoft.com/en-us/azure/azure-sql/database/business-continuity-high-availability-disaster-recover-hadr-overview?view=azuresql> Recovery method RTO RPO Geo-restore from geo-replicated backups 12 h 1 h Auto-failover groups 1 h 5 s Manual database failover 30 s 5 s

Auto fail over groups gives the lesser RPO&RTO and less administrative effort than manual fail over.

**NEW QUESTION: 184**

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

You have an application that queries DB1 to generate a sales report.

You need to see the parameter values from the last time the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Enable Last\_Query\_Plan\_Stats in the master database
- B. Enable Lightweight\_Query\_Profiling in DB1
- C. Enable Last\_Query\_Plan\_Stats in DB1
- D. Enable Lightweight\_Query\_Profiling in the master database
- E. Enable PARAMETER\_SNIFFING in DB1

**Answer: (SHOW ANSWER)**

LAST\_QUERY\_PLAN\_STATS works under the lightweight query execution statistics profiling infrastructure, so we need to ensure both Last\_Query\_Plan\_Stats Lightweight\_Query\_Profiling is on at database scope (though Lightweight\_Query\_Profiling is default on at SQL2019 but it can be turned off as well).

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-plan-stats-transact-sql?view=sql-server-ver15>

**NEW QUESTION: 185**

Drag and Drop Question

You have an instance of SQL Server on Azure Virtual Machines named SQL1. SQL1 contains a database named DB1.

You need to enable Transparent Data Encryption (TDE) for DB1.

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.

The screenshot shows a Microsoft drag-and-drop question interface. On the left, under the heading "Objects", there is a list of six items: "a database encryption key in the master database", "a master key in the master database", "a certificate in DB1", "a certificate in the master database", "a master key in DB1", and "a database encryption key in DB1". On the right, under the heading "Answer Area", there are three numbered slots (1, 2, 3) for placing the objects in sequence. There are arrows for moving objects from the list to the answer area and a vertical arrow for ordering the objects in the answer area. A watermark "freedam.net" is visible across the interface.

**Answer:**

**Objects**

a database encryption key in the master database

a certificate in DB1

a master key in DB1

**Answer Area**

1


a master key in the master database

2

a certificate in the master database

3

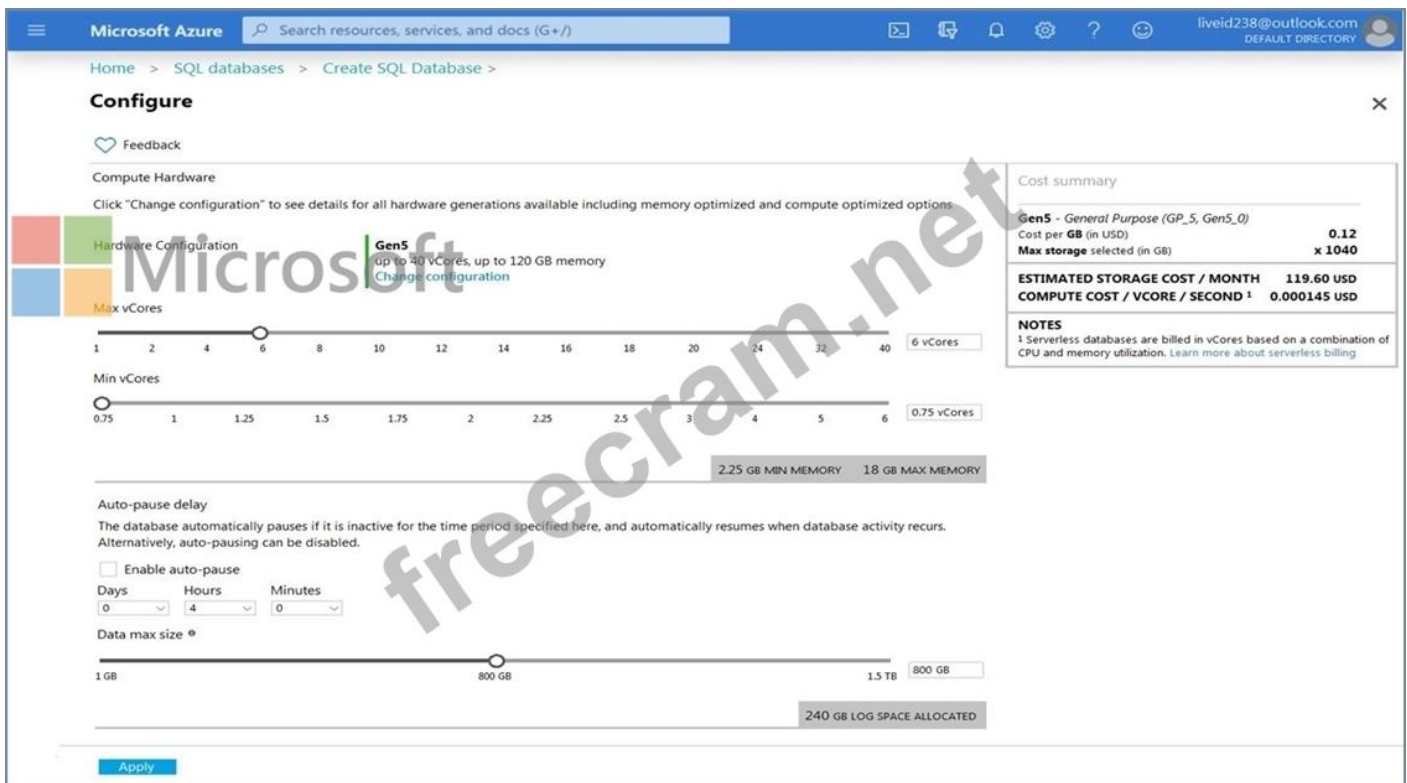
a database encryption key in DB1



**NEW QUESTION: 186**

Hotspot Question

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



The screenshot shows the 'Configure' page for an Azure SQL database. The 'Hardware Configuration' section is expanded, showing 'Gen5' with 'up to 40 vCores, up to 120 GB memory'. The 'Max vCores' slider is set to 6, and the 'Min vCores' slider is set to 0.75. The 'Auto-pause delay' is set to 4 hours. The 'Data max size' is set to 800 GB. The 'Cost summary' on the right shows a cost of 0.12 per GB and an estimated storage cost of 119.60 USD per month.

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

NOTE: Each correct selection is worth one point.

Answer Area



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:

**Answer Area**

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.

no extra time
up to 10 minutes
up to one minute

intermittent and unpredictable
regular and high
steady and low

Explanation:

Box 1: no extra time

Auto Pause is not checked in the exhibit.

Note: If Auto Pause is checked the correct answer is: up to one minute

Box 2: intermittent and unpredictable

Reference:


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

**NEW QUESTION: 187**


Hotspot Question

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.

**Answer Area**

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:**

**Answer Area**

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"/>

Explanation:

Box 1: No

Just one failure within the 5-minute interval.

Box 2: No

Just two failures within the 5-minute interval.

Box 3: No

Just two failures within the 5-minute interval.

Reference:

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

## NEW QUESTION: 188

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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.

Number of role assignments for this subscription: 15 / 2000

Name: Search by name or email | Type: Groups | 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)

<input type="checkbox"/>	Name	Type	Role	Scope
<input type="checkbox"/>	DBAGroup1	Group	Contributor	This resource
<input type="checkbox"/>	DBAGroup2	Group	SQL DB Contributor	This resource

Which audit log destination should you use to meet the monitoring requirements?

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

**Answer: (SHOW ANSWER)**

Scenario: Use a single dashboard to review security and audit data for all the PaaS databases. With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics.

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

Hubs.

Reference:

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

## **NEW QUESTION: 189**

Case Study 4 - A.Datum

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.

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](#))

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.

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-functions.md>

**NEW QUESTION: 190**

Case Study 1 - Litware, Inc

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.

## Drag and Drop Question

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

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

## Actions

## Answer Area

Modify the Azure SQL server administrator account.

Create database users.

Create a user in the master database.

Create an Azure AD administrator for the logical server.

Connect to the databases by using an Azure AD account.

Answer:

## Actions

## Answer Area

Modify the Azure SQL server administrator account.

Create an Azure AD administrator for the logical server.

Create database users.

Create a user in the master database.

Connect to the databases by using an Azure AD account.

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.

1. Create and populate Azure AD.
2. Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
3. Create an Azure Active Directory administrator. (Step 1)

4. Configure your client computers.
5. Create contained database users in your database mapped to Azure AD identities. (Step 2)
6. Connect to your database by using Azure AD identities. (Step 3)

Reference:

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

### NEW QUESTION: 191

Drag and Drop Question

You have an Azure virtual machine named VM1 that runs Red Hat Enterprise Linux (RHEL).

You need to install and configure Microsoft SQL Server 2022 Standard on VM1.

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.

**Answer:**

Explanation:

Reference:

<https://learn.microsoft.com/en-us/sql/linux/quickstart-install-connect-red-hat>

### NEW QUESTION: 192

You are monitoring an Azure Stream Analytics job.

You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero.

You need to ensure that the job can handle all the events.

What should you do?

- A. Remove any named consumer groups from the connection and use \$default.
- B. Change the compatibility level of the Stream Analytics job.

C. Create an additional output stream for the existing input stream.

D. Increase the number of streaming units (SUs).

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

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

Reference:

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

### NEW QUESTION: 193

Hotspot Question

You have an Azure subscription.

You need to deploy two virtual machines named VM1 and VM2 that will host Microsoft SQL Server instances. The solution must meet the following requirements:

- VM1 must provide the highest I/O throughput-to-vCore ratio.
- VM2 must provide the highest memory-to-vCore ratio.

Which series should you select for each virtual machine? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**  Microsoft

VM1:

	▼
DCadsv5-series	
Ebdsv5-series	
Mv2-series	

VM2:

	▼
DCadsv5-series	
Ebdsv5-series	
Mv2-series	

**Answer:**

/M1:   
DCadsv5-series  
EbdsV5-series  
Mv2-series

/M2:   
DCadsv5-series  
EbdsV5-series  
Mv2-series

Explanation:

Box 1: EbdsV5-series

VM1 must provide the highest I/O throughput-to-vCore ratio.

Consider the following VM series based on your SQL Server workloads:

High I/O throughput-to-vCore ratio: The Mbsv3 and MbdsV3-series VMs offer the highest throughput-to-vCore ratio of any VM series in any cloud. The EbdsV5-series VMs offer the second highest throughput-to-vCore ratio in Azure. If you don't know the I/O requirements for your SQL Server workload, the EbdsV5-series VMs is the one most likely to meet your needs.

Box 2: Mv2 series

VM2 must provide the highest memory-to-vCore ratio.

The Azure Virtual Machine Mv2 series offers the highest memory-to-vCore ratio among Azure VM sizes. These VMs are specifically designed for memory-intensive workloads like in-memory databases and large-scale analytics, providing up to 12TB of memory with a substantial number of vCPUs.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices-vm-size?view=azuresql>

<https://learn.microsoft.com/en-us/azure/virtual-machines/mv2-series>

**NEW QUESTION: 194**

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: C (LEAVE A REPLY)**

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: 195**

### **SIMULATION**

You need to ensure that when non-administrative users query the SalesLT.Customer table in db1, email addresses are obscured. For example, an email address of alice@contoso.com must appear as aXXX@XXXX.com.

You may need to use SQL Server Management Studio and the Azure portal.

### **Answer:**

SQL Database dynamic data masking with the Azure portal

In the Dynamic Data Masking pane in the portal, you see the recommended columns for your database. Select Add Mask for one or more columns, then select the appropriate masking function and select Save, to apply mask for these fields.

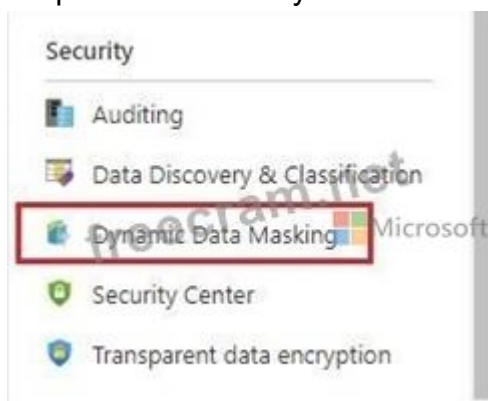
In details:

Enable dynamic data masking

Step 1: Launch the Azure portal at <https://portal.azure.com>.

Step 2: Go to your database resource in the Azure portal. [Select db1]

Step 3: Select the Dynamic Data Masking pane under the Security section.



Step 4: In the Dynamic Data Masking configuration page, you may see some database columns that the recommendations engine has flagged for masking. In order to accept the recommendations, just click Add Mask for one or more columns and a mask is created based on the default type for this column. You can change the masking function by clicking on the masking rule and editing the masking field format to a different format of your choice. Be sure to click Save to save your settings.

4.1 If possible locate the column with email addresses in the SalesLT.Customer table [If not seen

skip to step 5]

#### 4.2 Click Add Mask

Step 5: To add a mask for any column in your database, at the top of the Dynamic Data Masking configuration page, click Add Mask to open the Add Masking Rule configuration page.

[Skip steps 5-6 step if step 4 was successful]

**Add masking rule**

↑ Add ✕ Discard 🗑 Delete

Mask name  
Mask name

Select what to mask

Schema  
[Dropdown]

Table  
[Dropdown]

Column  
[Dropdown]

Select how to mask

Masking field format  
Default value (0, xxxx, 01-01-1900) [Dropdown]

Step 6: Select the Schema, Table and Column to define the designated field for masking.

Select the following

Table: SalesLT.Customer table

Column: [Choose the column with the email addresses]

Step 7: Select how to mask from the list of sensitive data masking categories.

Select Email (aXXX@XXXX.com)

## Add masking rule

 Add  Discard  Delete

Mask name

Mask name

Select what to mask

Schema

Table

Column

Select how to mask

Masking field format

Default value (0, xxxx, 01-01-1900) 

Default value (0, xxxx, 01-01-1900)

Credit card value (xxxx-xxxx-xxxx-1234)

Email (aXXX@XXXX.com)

Number (random number range)

Custom string (prefix [padding] suffix)



Step 8: Click Add in the data masking rule page to update the set of masking rules in the dynamic data masking policy.

Step 9: Type the SQL authenticated users or authenticated identities from Microsoft Entra ID (formerly Azure Active Directory) that should be excluded from masking, and have access to the unmasked sensitive data. [Skip] Step 10: Click Save in the data masking configuration page to save the new or updated masking policy.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-configure-portal>

### NEW QUESTION: 196

Drag and Drop Question

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 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:**

**Actions**

- Create a plan.
- Create a classifier function in db1.

**Answer Area**

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

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

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

DP-300 exam **questions have been updated** and **answers have been corrected** get the

**newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdumps.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

**Special Discount Code: freecram**)

**NEW QUESTION: 197**

You have an Azure subscription that contains two Azure SQL servers named Server1 and Server2. Server1 is in the East US Azure region and hosts a database named DB1. Server2 is in the West US Azure region.

You enable active geo-replication for DB1 and configure Server2 as the target server.

You need to test the availability of DB1 in the event of an Azure region outage. The solution must minimize administrative effort.

Which Azure Command-Line Interface (CLI) command should you include in the solution?

- A. az sql replica delete-link
- B. az sql db delete
- C. az sql db update
- D. az sql db rename

**Answer: (SHOW ANSWER)**

<https://learn.microsoft.com/en-us/cli/azure/sql/db?view=azure-cli-latest#az-sql-db-rename>

<https://learn.microsoft.com/en-us/cli/azure/sql/db/replica?view=azure-cli-latest#az-sql-db-replica-delete-link>

**NEW QUESTION: 198**

You have five on-premises servers that have Microsoft SQL Server 2022 Enterprise installed. Each server contains multiple database that store data for web apps and are backed up by using a third-party backup solution.

You plan to migrate the databases to Azure.

You need to recommend a solution to host the databases. The solution must meet the following requirements:

- Compute and storage resources must be shared across the databases.
- Costs must be minimized.

What should you include in the recommendation?

- A. Azure SQL Managed Instance
- B. Azure SQL Database
- C. an Azure SQL Database elastic pool
- D. SQL Server on Azure Virtual Machines

**Answer: (SHOW ANSWER)**

SQL Server VM alternative

Your business might have requirements that make SQL Server on Azure Virtual Machines a more suitable target than Azure SQL Database.

Reference:

<https://learn.microsoft.com/en-us/data-migration/sql-server/database/overview>

**Valid DP-300 Dumps** shared by ExamDiscuss.com for Helping Passing DP-300 Exam!  
ExamDiscuss.com now offer the **newest DP-300 exam dumps**, the ExamDiscuss.com  
DP-300 exam **questions have been updated** and **answers have been corrected** get the  
**newest** ExamDiscuss.com DP-300 dumps with Test Engine here:

<https://www.examdiscuss.com/Microsoft/exam/DP-300/premium/> (368 Q&As Dumps, **35%OFF**

**Special Discount Code: [freecram](#)**)