

Microsoft.DP-300.vJan-2024.by.Ricky.112q

Number: DP-300
Passing Score: 800
Time Limit: 120
File Version: 21.0

Exam Code: DP-300
Exam Name: Administering Relational Databases on Microsoft Azure



01 - Plan and Implement Data Platform Resources

Case study

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

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

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

To start the case study

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

When you are ready to answer a question, click the Question button to return to the question.

Overview

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.



QUESTION 1

What should you use to migrate the PostgreSQL database?

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

Correct Answer: C

Section:

Explanation:

Reference:

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

02 - Plan and Implement Data Platform Resources

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

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

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

To start the case study

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

Overview

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

Physical Locations

Existing Environment

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

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

Network Environment

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

Identity Environment

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

Database Environment

The sales department has the following database workload:

An on-premises server named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases. A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1.

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

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

Current Problems

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

Requirements

Planned Changes

Litware plans to implement the following changes:

QUESTION 1

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Azure Database Migration Service pricing tier:

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

Required Azure resource:

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

Answer Area:

Answer Area



Azure Database Migration Service pricing tier:

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

Required Azure resource:

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

Section:

Explanation:

Azure Database Migration service

Box 1: Premium 4-VCore

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

Minimize downtime during the migration of the SERVER1 databases.

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

Incorrect Answers:

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

Box 2: A VPN gateway

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

Reference:

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

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

QUESTION 2

HOTSPOT

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Purchasing model:

- Azure virtual machine reserved instances
- DTU
- vCore

Deployment option:

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

Answer Area:

Answer Area

Purchasing model: ▼

- Azure virtual machine reserved instances
- DTU
- vCore

Deployment option: ▼

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

Section:

Explanation:

Box 1: DTU

Scenario:

The 30 new databases must scale automatically.

Once all requirements are met, minimize costs whenever possible.

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model. In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.

Box 2: An Azure SQL database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

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

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



03 - Plan and Implement Data Platform Resources

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

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

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

To start the case study

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

When you are ready to answer a question, click the Question button to return to the question.

Overview

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:

QUESTION 1

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



Correct Answer: C

Section:

Explanation:

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>

QUESTION 2

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

Correct Answer: B

Section:

Explanation:

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>

QUESTION 3

HOTSPOT

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.

Hot Area:

Answer Area

Table type to store retail store data:

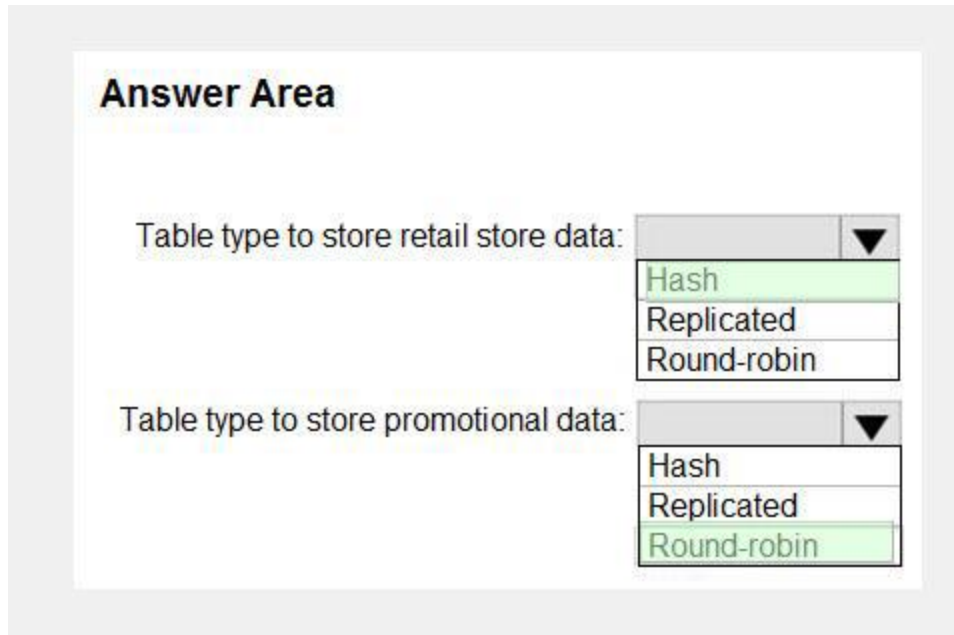
Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Answer Area:





Section:

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/>

04 - Plan and Implement Data Platform Resources

QUESTION 1

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

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

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

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

Correct Answer: D

Section:

QUESTION 2

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

Correct Answer: A, C, E

Section:

Explanation:

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

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

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

Reference:

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

QUESTION 3

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

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)



FactSales has 6 billion rows and is loaded nightly by using a batch process.

Which type of compression provides the greatest space reduction for the database?

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

Correct Answer: D

Section:

Explanation:

Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.

Note: Columnstore — The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.

Incorrect Answers:

B: Rowstore — The rowstore index is the traditional style that has been around since the initial release of SQL Server.

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database.

Reference:

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

QUESTION 4

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

Correct Answer: B

Section:

Explanation:

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

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

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>

QUESTION 5

HOTSPOT

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

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

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

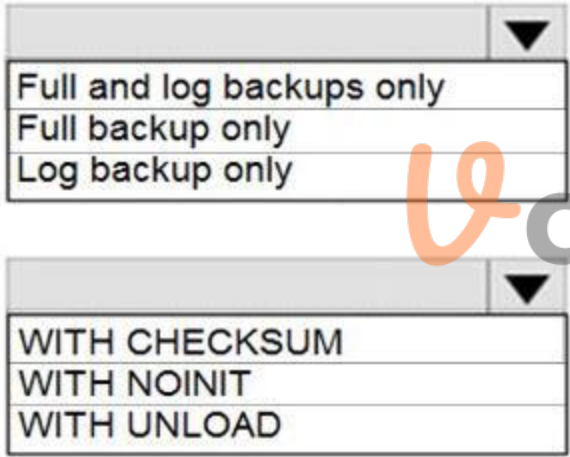
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Backup type:

Backup option:



Answer Area:

Answer Area

Backup type:

Backup option:

Section:

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>

QUESTION 6

DRAG DROP

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.

Select and Place:

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.



Correct Answer:

Actions

Answer Area

- Change the Active Directory Admin on TestServer1
-
-
- Add IP addresses to the firewall
-
-



- From the Azure portal, export the Azure Resource Manager templates
- Change the server name and related variables in the templates
- From the Azure portal, deploy the templates.
- From the database project, deploy the database schema and permissions



Section:
Explanation:

QUESTION 7
DRAG DROP

You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database. A recent failure causes the master database to be lost. You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.

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

Select and Place:

Actions		Answer Area
Add a certificate to an Azure key vault		
Enable Transparent Data Encryption (TDE)		
Encrypt a copy of the master key by using the service master key	➤	⬆
Turn on the TRUSTWORTHY property and the CLR property	➤	⬇
Attach the SSISDB database		
Open the master key for the SSISDB database		

Correct Answer:



Actions		Answer Area
Add a certificate to an Azure key vault		Attach the SSISDB database
Enable Transparent Data Encryption (TDE)		Turn on the TRUSTWORTHY property and the CLR property
	➤	Open the master key for the SSISDB database
	➤	Encrypt a copy of the master key by using the service master key

Section:

Explanation:

Step 1: Attach the SSISDB database

Step 2: Turn on the TRUSTWORTHY property and the CLR property

If you are restoring the SSISDB database to an SQL Server instance where the SSISDB catalog was never created, enable common language runtime (clr)

Step 3: Open the master key for the SSISDB database

Restore the master key by this method if you have the original password that was used to create SSISDB.

open master key decryption by password = 'LS1Setup!' --'Password used when creating SSISDB'

Alter Master Key Add encryption by Service Master Key

Step 4: Encrypt a copy of the mater key by using the service master key

Reference:

<https://docs.microsoft.com/en-us/sql/integration-services/backup-restore-and-move-the-ssis-catalog>

QUESTION 8

You are designing a streaming data solution that will ingest variable volumes of data.

You need to ensure that you can change the partition count after creation.

Which service should you use to ingest the data?

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

Correct Answer: D

Section:

Explanation:

The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created. Incorrect Answers:

A: For Azure Event standard hubs, the partition count isn't changeable, so you should consider long-term scale when setting partition count.

Reference:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

QUESTION 9

You have an Azure Synapse Analytics Apache Spark pool named Pool1.

You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file. You need to load the files into the tables. The solution must maintain the source data types.

What should you do?

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

Correct Answer: B

Section:

Explanation:

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools. Serverless SQL pool enables you to query data in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.

To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.

The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files>

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

QUESTION 10

You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables. Which distribution type should you recommend to minimize data movement?

- A. HASH
- B. REPLICATE
- C. ROUND_ROBIN

Correct Answer: B

Section:

Explanation:

A replicated table has a full copy of the table available on every Compute node. Queries run fast on replicated tables since joins on replicated tables don't require data movement. Replication requires extra storage, though, and isn't practical for large tables.

Incorrect Answers:

C: A round-robin distributed table distributes table rows evenly across all distributions. The assignment of rows to distributions is random. Unlike hash-distributed tables, rows with equal values are not guaranteed to be assigned to the same distribution.

As a result, the system sometimes needs to invoke a data movement operation to better organize your data before it can resolve a query.

Reference:

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

QUESTION 11

You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1. You plan to create a database named DB1 in Pool1.

You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool. Which format should you use for the tables in DB1?

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



Correct Answer: C

Section:

Explanation:

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools. For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.

Reference:

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

QUESTION 12

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

Send the output to an Azure Synapse.

Identify spikes and dips in time series data.

Minimize development and configuration effort.

Which should you include in the solution?

- A. Azure SQL Database
- B. Azure Databricks
- C. Azure Stream Analytics

Correct Answer: C

Section:

Explanation:

Anomalies can be identified by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics

Reference:

<https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/> <https://docs.microsoft.com/en-us/azure/stream-analytics/azure-synapse-analytics-output>

QUESTION 13

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL. Which switch should you use to switch between languages?

- A. \[<language>]
- B. %<language>
- C. \[<language>]
- D. @<language>

Correct Answer: B

Section:

Explanation:

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

Reference:

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

QUESTION 14

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

Correct Answer: A

Section:

Explanation:

Complete mode: You can use Structured Streaming to replace the entire table with every batch.

Incorrect Answers:

B: By default, streams run in append mode, which adds new records to the table.

Reference:

<https://docs.databricks.com/delta/delta-streaming.html>

QUESTION 15

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

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 16

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

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

Correct Answer: C

Section:

Explanation:

You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.

Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.

Reference:

<https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html>

QUESTION 17

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

Correct Answer: B

Section:

Explanation:

Instead use a serverless SQL pool to create an external table with the extra column.

Reference:

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

QUESTION 18

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

After you answer a question-in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1. You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1. Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column. Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section:

Explanation:

In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:

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

QUESTION 19

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

Correct Answer: B

Section:

Explanation:

Instead use a serverless SQL pool to create an external table with the extra column.

Note: In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:

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

QUESTION 20

HOTSPOT

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

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

You need to create a data source in Pool1 that you can reference when you create the external table.

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

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


|       |
|-------|
| blob  |
| dfs   |
| table |

  
        .core.windows.net',  
  


|                     |
|---------------------|
| PUSHDOWN = ON       |
| TYPE = BLOB_STORAGE |
| TYPE = HADOOP       |

  
    )
```

Answer Area:

Answer Area

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


|       |
|-------|
| blob  |
| dfs   |
| table |

  
        .core.windows.net',  
  


|                     |
|---------------------|
| PUSHDOWN = ON       |
| TYPE = BLOB_STORAGE |
| TYPE = HADOOP       |

  
    )
```

Section:

Explanation:

Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2

 Vdumps

```
CREATE EXTERNAL DATA SOURCE YellowTaxi
WITH ( LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/',
TYPE = HADOOP)
Box 2: HADOOP
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables
```

QUESTION 21

HOTSPOT

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

ProductID

ItemPrice

LineTotal

Quantity

StoreID

Minute

Month

Hour

Year

Day

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

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

```
df.write
```

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

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

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

Answer Area:

Answer Area

```
df.write
```

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

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

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

Section:

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>

QUESTION 22

Hot Area:



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 Area:

 Vdumps

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
```

Vdumps

Section:

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>

QUESTION 23

Hot Area:

Answer Area

EventCategory: ▼
DimChannel
DimDate
DimEvent
FactEvents

ChannelGrouping: ▼
DimChannel
DimDate
DimEvent
FactEvents

TotalEvents: ▼
DimChannel
DimDate
DimEvent
FactEvents

Answer Area:



Answer Area

EventCategory: ▼

- DimChannel
- DimDate
- DimEvent
- FactEvents

ChannelGrouping: ▼

- DimChannel
- DimDate
- DimEvent
- FactEvents

TotalEvents: ▼

- DimChannel
- DimDate
- DimEvent
- FactEvents



Section:

Explanation:

Box 1: FactEvents

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

Box 2: DimChannel

Dimension tables describe business entities – the things you model. Entities can include products, people, places, and concepts including time itself. The most consistent table you'll find in a star schema is a date dimension table. A dimension table contains a key column (or columns) that acts as a unique identifier, and descriptive columns.

Box 3: DimEvent

Reference:

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

QUESTION 24

DRAG DROP

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.

Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted.

You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Values

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

Answer Area

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

Correct Answer:

Values

- CustomerKey
- ROUND_ROBIN
- REPLICATE
- SalesOrderNumber

Answer Area

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



Section:

Explanation:

Box 1: HASH

Box 2: OrderDateKey

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

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management. For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:

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

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

QUESTION 25

DRAG DROP

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.

Select and Place:

Actions

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 Area



Correct 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.

Section:

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>

QUESTION 26

HOTSPOT

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

Git repository

Git repository information associated with your data factory. [CI/CD best practices](#)

Setting Disconnect

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



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
 NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

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

<input type="text" value=""/>	▼
/	
adf_publish	
main	
Parameterization template	
<input type="text" value=""/>	▼
/contososales	
/dwh_batchetl/adf_publish/contososales	
/main	

Answer Area:

Answer Area

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

- /
- adf_publish
- main
- Parameterization template

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

- /contososales
- /dwh_batchetl/adf_publish/contososales
- /main

Section:

Explanation:

Box 1: adf_publish

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

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

Box 2: /dwh_barchlet/ adf_publish/contososales

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

Reference:

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



QUESTION 27

HOTSPOT

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.

Hot Area:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

- | |
|---|
| <input type="text"/> |
| CSV files in the tripdata_2020 subfolder |
| files that have file names beginning with "tripdata_2020" |
| CSV files that have file names containing "tripdata_202" |
| CSV files that have file named beginning with "tripdata_2020" |

The query assumes that the first row in a CSV file is **[answer choice]** row.

- | |
|----------------------|
| <input type="text"/> |
| a header |
| a data |
| an empty |

Answer Area:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

▼
CSV files in the tripdata_2020 subfolder
files that have files names beginning with "tripdata_2020"
CSV files that have file names containing "tripdata_202"
CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.

▼
a header
a data
an empty

Section:

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>



QUESTION 28

HOTSPOT

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

Configure

Feedback

Compute Hardware

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

Hardware Configuration

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

Max vCores



Min vCores



2.25 GB MIN MEMORY 18 GB MAX MEMORY

Auto-pause delay

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

Enable auto-pause

Days: 0 Hours: 4 Minutes: 0

Data max size



240 GB LOG SPACE ALLOCATED

Apply

Cost summary

Gen5 - General Purpose (GP_5, Gen5_0)
Cost per GB (in USD) **0.12**
Max storage selected (in GB) **x 1040**

ESTIMATED STORAGE COST / MONTH 119.60 USD
COMPUTE COST / VCORE / SECOND ¹ 0.000145 USD

NOTES

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

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.

Hot Area:

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 Area:

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

Section:

Explanation:

Reference:

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

QUESTION 29

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

The web app must be hosted on an Azure virtual network.

The Azure SQL database must be assigned a private IP address.

The Azure SQL database must allow connections only from the virtual network.
You need to recommend a solution that meets the requirements.
What should you include in the recommendation?

- A. Azure Private Link
- B. a network security group (NSG)
- C. a database-level firewall
- D. a server-level firewall

Correct Answer: A

Section:

QUESTION 30

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

Correct Answer: A

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vcare-single-databases#business-critical---provisioned-compute---gen4>

QUESTION 31

HOTSPOT

You have an Azure subscription.

You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template.

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

```
"resources": [  
  ...  
  "type":   
    Microsoft.Sql/servers  
    Microsoft.Sql/servers/databases  
    Microsoft.Sql/managedInstances  
  "name": "[parameters('targetName')]",  
  "location": "[parameters('location')]",  
  "sku": {  
    "name": "[parameters('skuName')]"  
  }  
  ...  
  "dependsOn": [  
      
    "[parameters('targetName')]",  
    "[parameters('virtualNetworkName')]",  
    "[variables('networkSecurityGroupName')]",  
  ],  
  "properties": {  
    "administratorLogin": "[parameters('administratorLogin')]",  
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",  
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'),  
    parameters('virtualNetworkName'), parameters('subnetName'))]",  
    "storageSizeInGB": "[parameters('storageSizeInGB')]", "vCores": "[parameters('vCores')]",  
    "licenseType": "[parameters('licenseType')]"  
  }  
  ...  
]
```

Answer Area:

Answer Area

```
"resources": [  
  ...  
  "type":   
    Microsoft.Sql/servers  
    Microsoft.Sql/servers/databases  
    Microsoft.Sql/managedInstances  
  "name": "[parameters('targetName')]",  
  "location": "[parameters('location')]",  
  "sku": {  
    "name": "[parameters('skuName')]"  
  }  
  ...  
  "dependsOn": [  
      
    "[parameters('targetName')]",  
    "[parameters('virtualNetworkName')]",  
    "[variables('networkSecurityGroupName')]",  
  ],  
  "properties": {  
    "administratorLogin": "[parameters('administratorLogin')]",  
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",  
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'),  
    parameters('virtualNetworkName'), parameters('subnetName'))]",  
    "storageSizeInGB": "[parameters('storageSizeInGB')]", "vCores": "[parameters('vCores')]",  
    "licenseType": "[parameters('licenseType')]"  
  }  
  ...  
]
```

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/create-template-quickstart?tabs=azure-powershell>

QUESTION 32

HOTSPOT

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.

Hot Area:

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 Area:

Answer Area

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"/>

Section:

Explanation:

Reference:

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

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

QUESTION 33

HOTSPOT

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.

Hot Area:

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 Area:

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

Section:

Explanation:

Reference:

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

01 - Implement a Secure Environment

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

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

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

To start the case study

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

When you are ready to answer a question, click the Question button to return to the question.

Overview

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

Physical Locations

Existing Environment

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

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

Network Environment

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

Identity Environment

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

Database Environment

The sales department has the following database workload:

An on-premises server named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases. A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1.

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

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

Current Problems

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

Requirements

Planned Changes

Litware plans to implement the following changes:

Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB. Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data. Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1. Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

Maintenance tasks must be automated.

The 30 new databases must scale automatically.

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

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

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

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

Store encryption keys in Azure Key Vault.

Retain backups of the PII data for two months.

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

Use the principle of least privilege whenever possible.

Authenticate database users by using Active Directory credentials.

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

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

Business Requirements

Litware identifies the following business requirements:

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

Minimize downtime during the migration of the SERVER1 databases.

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

Once all requirements are met, minimize costs whenever possible.

QUESTION 1

DRAG DROP

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.

Select and Place:

Actions

Create a user in the master database

Modify the Azure SQL server administrator account

Create contained database users

Create an Azure AD administrator for the logical server

Connect to the databases by using an Azure AD account

Enable the contained database authentication option

Answer Area



Correct Answer:

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

Section:

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. Connect to the databases using an Azure AD account (the Administrator account that was configured in the previous step). (Step 2)
5. Create contained database users in your database mapped to Azure AD identities. (Step 3)

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-configure?tabs=azure-powershell>

QUESTION 2

DRAG DROP

You create all of the tables and views for ResearchDB1.

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

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

Select and Place:

Actions

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

Answer Area



Correct 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.

Section:

Explanation:

Reference:

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

02 - Implement a Secure Environment

Case study

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

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

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

return to this section.

To start the case study

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

When you are ready to answer a question, click the Question button to return to the question.

Overview

Existing Environment

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

Active Directory

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

Database Environment

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

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

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

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

Current Business Model

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

Requirements

Planned Changes

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

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.

Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months.

Start onboarding customers to the new PaaS solution within six months.

Business Goals

Contoso identifies the following business requirements:

Use built-in Azure features whenever possible.

Minimize development effort whenever possible.

Minimize the compute costs of the PaaS solutions.

Provide all the customers with their own copy of the database by using the PaaS solution.

Provide the customers with different table and row access based on the customer's service agreement. In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover. Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso.

Technical Requirements

Contoso identifies the following technical requirements:

Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure AD of Contoso to minimize guest users.

All customers must have their own resource group, Azure SQL server, and Azure SQL database. The deployment of resources for each customer must be done in a consistent fashion. Users must be able to review the queries issued against the PaaS databases and identify any new objects created. Downtime during the PostgreSQL database migration must be minimized.

Monitoring Requirements

Contoso identifies the following monitoring requirements:

Notify administrators when a PaaS database has a higher than average CPU usage.

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

Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.


Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues automatically whenever possible.

PaaS Prototype

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

Show data for last: 1 hour 24 hours 7 days Aggregation type: Max

Compute utilization



Role Assignments

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

+ Add Edit columns Refresh Remove Got feedback?

Check access **Role assignments** Deny assignments Classic administrators Roles

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

Number of role assignments for this subscription 15 2000

Name Type Groups Role 2 selected Scope All scopes

Group by Role

i 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"/>	DB DBAGroup1	Group	Contributor	This resource
<input type="checkbox"/>	DB DBAGroup2	Group	SQL DB Contributor	This resource

QUESTION 1

HOTSPOT

You are evaluating the role assignments.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Answer Area:

Answer Area

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

Section:

Explanation:

Box 1: Yes

DBAGroup1 is member of the Contributor role.

The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.

Box 2: No

Box 3: Yes

DBAGroup2 is member of the SQL DB Contributor role.

The SQL DB Contributor role lets you manage SQL databases, but not access to them. Also, you can't manage their security-related policies or their parent SQL servers. As a member of this role you can create and manage SQL databases.

Reference:

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

QUESTION 2

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.

Correct Answer: B

Section:

QUESTION 3

You are evaluating the business goals.

Which feature should you use to provide customers with the required level of access based on their service agreement?

- A. dynamic data masking
- B. Conditional Access in Azure
- C. service principals
- D. row-level security (RLS)

Correct Answer: D

Section:

Explanation:

Reference:

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

03 - Implement a Secure Environment

QUESTION 1

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).

Correct Answer: A, C, D

Section:

Explanation:

C: Advanced Data Security (ADS) is a unified package for advanced SQL security capabilities. ADS is available for Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics. It includes functionality for discovering and classifying sensitive data

D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes that have been added to the SQL Server database engine. This metadata can then be used for advanced, sensitivity-based auditing and protection scenarios.

A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called data_sensitivity_information. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	██████████7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	██████████7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	██████████7.125	Microsoft SQL Server Management Studio - Query	41	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview>

QUESTION 2

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.

Minimize the possibility of VM1 exfiltrating data stored in SqlDb1.

What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link
- D. an ExpressRoute gateway

Correct Answer: C

Section:

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network. Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

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

QUESTION 3

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.

Correct Answer: B, E

Section:

Explanation:

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



to different availability zones in the same region. To eliminate a singlepoint of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zoneredundant configuration in the Premium or Business Critical service tiers does not create additionaldatabase redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration,you can make your Premium or Business Critical databases resilient to a much larger set of failures,including catastrophic datacenter outages, without any changes to the application logic. You can alsoconvert any existing Premium or Business Critical databases or pools to the zone redundantconfiguration.

QUESTION 4

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.

Restrict network connectivity to SqlSrv1.

What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link
- D. an ExpressRoute gateway

Correct Answer: C

Section:

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

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

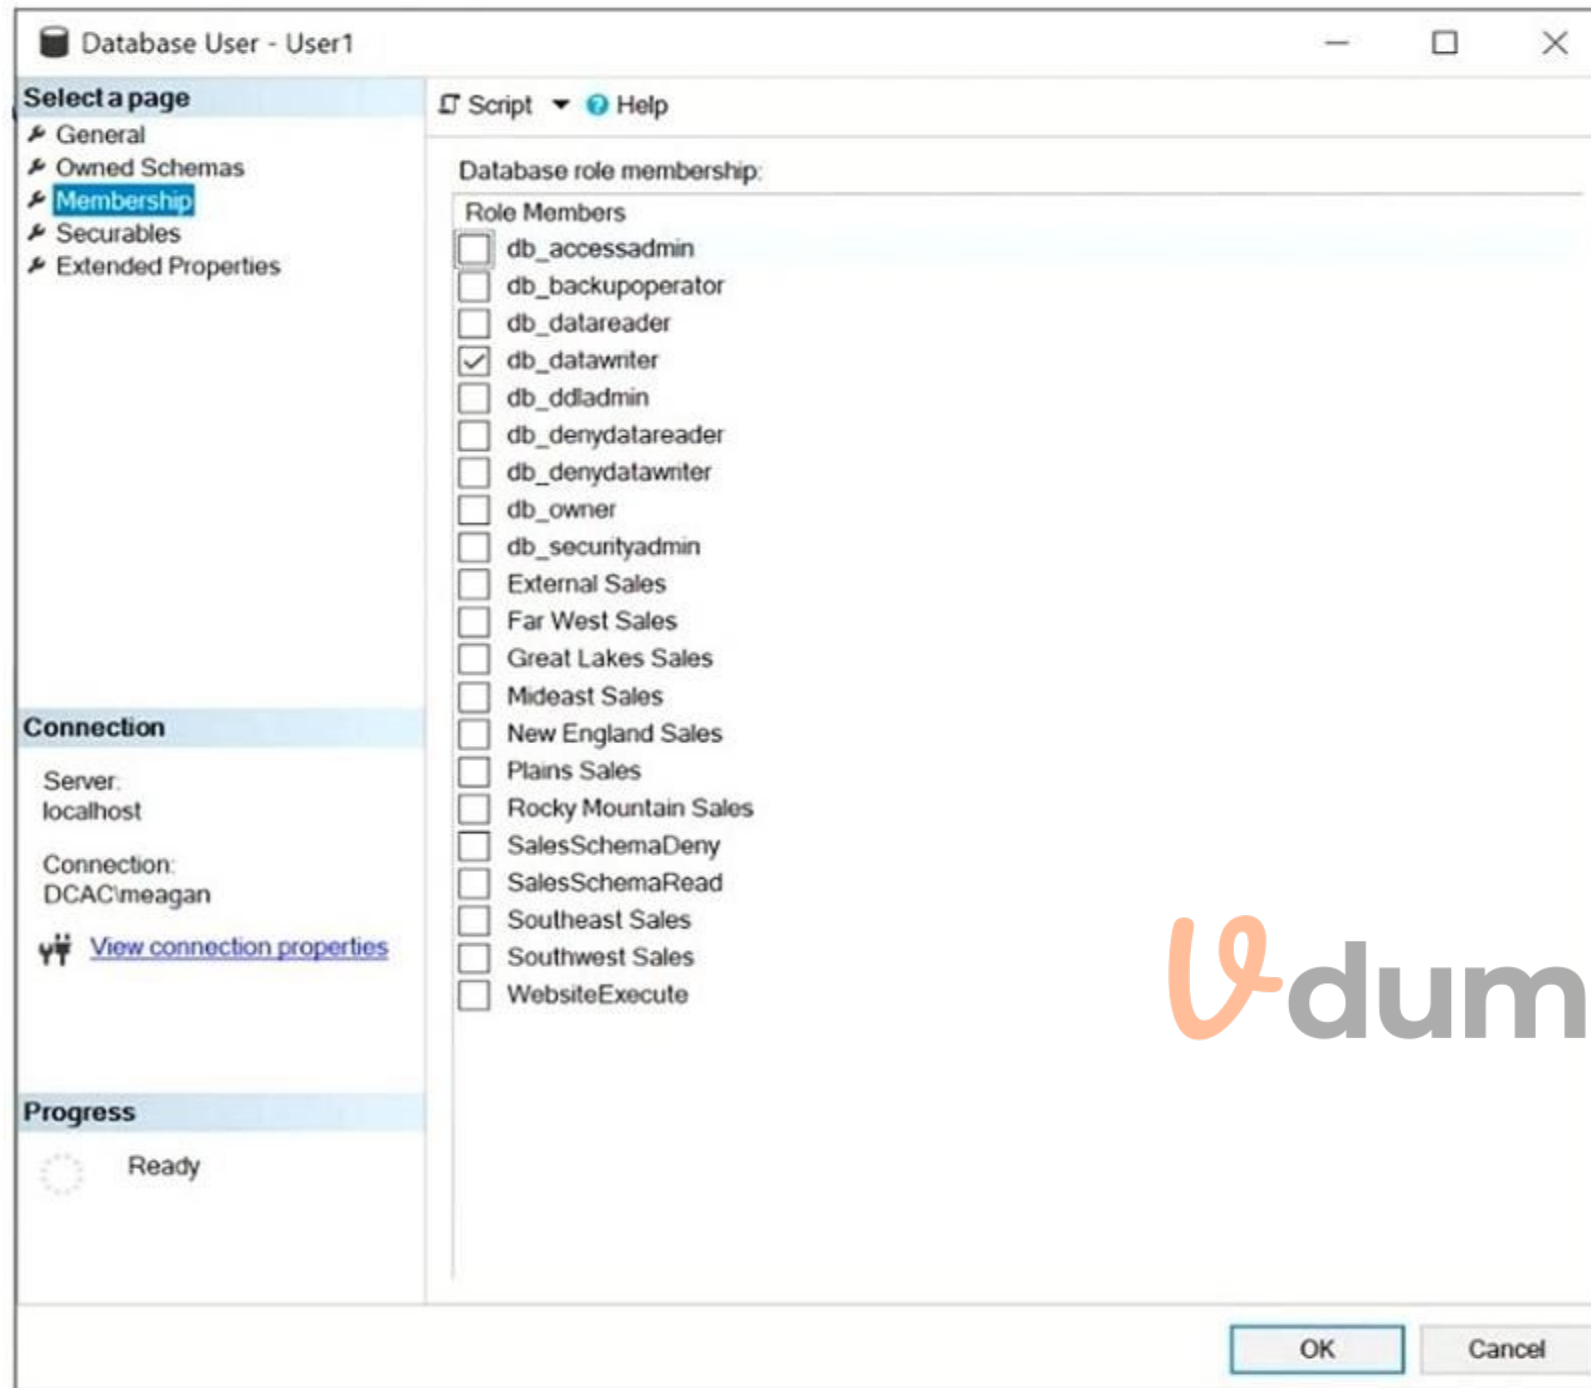


QUESTION 5

HOTSPOT

You have a Microsoft SQL Server database named DB1 that contains a table named Table1.

The database role membership for a user named User1 is shown in the following exhibit.



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

User1 can [answer choice].

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

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

	▼
db_datareader	
db_ddladmin	
db_denydatareader	
db_denydatawriter	

Answer Area:

Answer Area

User1 can [answer choice].

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

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

	▼
db_datareader	
db_ddladmin	
db_denydatareader	
db_denydatawriter	

Section:

Explanation:

Box 1: delete a row from Table1

Members of the db_datawriter fixed database role can add, delete, or change data in all user tables.

Box 2: db_datareader

Members of the db_datareader fixed database role can read all data from all user tables.

Reference:

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

QUESTION 6

DRAG DROP

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1.

The only user who was created is the server administrator.

You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.

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

Select and Place:

Actions

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.

Answer Area



Correct Answer:

Actions

Connect to DB1 by using the Active Directory admin account.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.

Answer Area

Set the Active Directory Admin for AzSQL1.

Connect to DB1 by using the server administrator account.



Section:

Explanation:

Step 1: Set up the Active Directory Admin for AzSQL1.

Step 2: Connect to DB1 by using the server administrator.

Sign into your managed instance with an Azure AD login granted with the sysadmin role.

Step 3: Create a user by using the FROM EXTERNAL PROVIDER clause.

FROM EXTERNAL PROVIDER is available for creating server-level Azure AD logins in SQL Database managed instance. Azure AD logins allow database-level Azure AD principals to be mapped to server-level Azure AD logins. To create an Azure AD user from an Azure AD login use the following syntax:

```
CREATE USER [AAD_principal] FROM LOGIN [Azure AD login]
```

Reference:

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

QUESTION 7

HOTSPOT

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

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

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

- The first six numerals of each customer's phone number must be masked.
- The first six numerals of each customer's phone number must be masked.
- The last four digits of each customer's phone number must be visible.

Hyphens must be preserved and displayed.

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

Hot Area:

Answer Area



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 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	



Section:

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>

QUESTION 8

DRAG DROP

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.

Select and Place:

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.



Correct Answer:

Actions

Answer Area

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.

Create a column encryption key.



Encrypt the Salary column by using the randomized encryption type.



Section:

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>

QUESTION 9

HOTSPOT

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.

Hot Area:

Answer Area

Encryption key: ▼

Column encryption key
Database encryption key
Service master key

Encryption type: ▼

Deterministic
Randomized
Transparent Data Encryption (TDE)



Answer Area:

Answer Area

Encryption key: ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type: ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

Section:

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>

QUESTION 10

DRAG DROP

You have an Azure SQL Database instance named DatabaseA on a server named Server1.

You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.

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

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

Select and Place:

Actions

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

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

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

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

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

Answer Area



Correct Answer:

Actions

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

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

Answer Area

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

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



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



Section:

Explanation:

Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'

Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:

```
CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
```

You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command. Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:

```
CREATE USER readonlyuser FROM LOGIN readonlylogin;
```

Step 3: On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]

Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to the database via the

db_datareader role.

```
EXEC sp_addrolemember 'db_datareader', 'readonlyuser';
```

Reference:

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

QUESTION 11

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

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

Correct Answer: C

Section:

Explanation:

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

What resources the client may access.

What permissions they have to those resources.

How long the SAS is valid.

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

Shared Key authorization

Shared access signature (SAS) authorization

Role-based access control (Azure RBAC)

Access control lists (ACL) Data Lake Storage Gen2 supports the following authorization mechanisms:

Shared Key authorization

Shared access signature (SAS) authorization

Role-based access control (Azure RBAC)

Access control lists (ACL)

Reference:

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

QUESTION 12

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

Correct Answer: B

Section:

Explanation:

Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users. The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card. Example:

XXXX-XXXX-XXXX-1234



QUESTION 13

You have a data warehouse in Azure Synapse Analytics.

You need to ensure that the data in the data warehouse is encrypted at rest.

What should you enable?

- A. Transparent Data Encryption (TDE)
- B. Advanced Data Security for this database
- C. Always Encrypted for all columns
- D. Secure transfer required

Correct Answer: A

Section:

Explanation:

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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview>

QUESTION 14

You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies. You need to ensure that users from each company can view only the data of their respective company.

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

- A. a column encryption key
- B. asymmetric keys
- C. a function
- D. a custom role-based access control (RBAC) role
- E. a security policy



Correct Answer: D, E

Section:

Explanation:

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes. Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac> <https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline>

QUESTION 15

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.

Correct Answer: B

Section:

Explanation:

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

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

Incorrect Answers:

A, D: Should enable Soft Delete and Do Not Purge on Azure Key Vault.

Using customer-managed keys with Data Factory requires two properties to be set on the Key Vault, Soft Delete and Do Not Purge. These properties can be enabled using either PowerShell or Azure CLI on a new or existing key vault.

Reference:

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

QUESTION 16

HOTSPOT

You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.

You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.

Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

MFA:


	▼
Azure AD authentication	
Microsoft SQL Server authentication	
Passwordless authentication	
Windows authentication	

Database-level authentication:

	▼
Application roles	
Contained database users	
Database roles	
Microsoft SQL Server logins	

Vdumps

Answer Area:



Section:

Explanation:

Box 1: Azure AD authentication

Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.

Box 2: Contained database users

Azure Active Directory Uses contained database users to authenticate identities at the database level.

Incorrect:

SQL authentication: To connect to dedicated SQL pool (formerly SQL DW), you must provide the following information:

Fully qualified servername

Specify SQL authentication

Username

Password

Default database (optional)

Reference:

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

QUESTION 17

You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.

You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication.

You need to ensure that App1 can access DB2. The solution must meet the following requirements:

App1 must be able to view only DB2.

Administrative effort must be minimized.

What should you create?

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

Correct Answer: D

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-database-portable?view=sql-server-ver15>

QUESTION 18

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

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1.

User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.

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

What should you do?

- A. Create User1 in the master database.
- B. Assign User1 the db_datareader role for the master database.
- C. Assign User1 the db_datareader role for the databases that User1 creates.
- D. Grant SELECT on sys.databases to public in the master database.

Correct Answer: A

Section:

Explanation:

Reference:

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

Exam H

QUESTION 1

You have an Azure SQL managed instance.

You need to enable SQL Agent Job email notifications.

What should you do?

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

Correct Answer: C

Section:

Explanation:

Reference: <https://techcommunity.microsoft.com/t5/azure-sql-blog/sending-emails-in-azure-sql-managed-instance/ba-p/386235>

QUESTION 2

You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases.

You need to update the column and index statistics for the databases.

What should you use?

- A. an Azure Automation runbook
- B. a SQL Agent job
- C. Azure SQL Analytics



D. automatic tuning in Azure SQL Database

Correct Answer: A

Section:

Explanation:

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

QUESTION 3

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.

Correct Answer: B, C

Section:

Explanation:

QUESTION 4

HOTSPOT

You have an Azure subscription.

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

How should you complete the template? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

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

Answer Area:

Answer Area

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

Section:

Explanation:

Reference:

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

QUESTION 5

HOTSPOT

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.

Hot Area:

Answer Area

EXECUTE AS

	▼
--	---

 = 'user1@contoso.com'

CALLER

LOGIN

OWNER

USER

GO

SELECT SUSER_SNAME ()

REVERT

REVOKE

ROLLBACK

GO

 **vdumps**

Answer Area:

Answer Area

EXECUTE AS

	▼
CALLER	
LOGIN	
OWNER	
USER	

 = 'user1@contoso.com'

	▼
CALLER	
LOGIN	
OWNER	
USER	

GO

SELECT SUSER_SNAME ()

	▼
REVERT	
REVOKE	
ROLLBACK	

GO

Vdumps

Section:

Explanation:

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>

QUESTION 6

DRAG DROP

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1.

You need to encrypt the data in Col1.

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

Select and Place:

Actions	Answer Area
Create a database master key.	
Create a column master key.	
Open the symmetric key.	
Create a certificate.	
Update Col1.	
Create a symmetric key.	

Correct Answer:

Actions	Answer Area
	Create a database master key.
Create a column master key.	Create a certificate.
	Create a symmetric key.
	Open the symmetric key.
Update Col1.	

Section:

Explanation:

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

QUESTION 7

HOTSPOT

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

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

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

How should you complete the Transact-SQL query? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

SELECT

	▼
is_user_process	
wait_time	
wait_type	

SUM(wait_time) AS total_wait_time_ms

FROM sys.

	▼
dm_exec_query_stats	
dm_exec_requests	
query_store_query	

```
JOIN sys.dm_exec_sessions AS dmvs2
  ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

Answer Area:

Answer Area

SELECT

is_user_process
wait_time
wait_type

SUM(wait_time) AS total_wait_time_ms

FROM sys.

dm_exec_query_stats
dm_exec_requests
query_store_query

```
JOIN sys.dm_exec_sessions AS dm2
  ON dm1.session_id = dm2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

Section:

Explanation:

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

QUESTION 8

HOTSPOT

You configure a long-term retention policy for an Azure SQL database as shown in the exhibit. (Click the Exhibit tab.)

Configure policies ✕

SQL server

Point in Time Restore Configuration

Configure PiTR backup retention ▼ Days

Long-term Retention Configurations

Weekly LTR Backups ⓘ

How long would you like weekly backups to be kept?

6 Week(s) ▼

Monthly LTR Backups ⓘ

How long would you like the first backup of each month to be kept?

12 Month(s) ▼

Yearly LTR Backups ⓘ

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

Week 2 ▼

How long would you like this annual backup to be kept?

10 Year(s) ▼

Vdumps

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.

Hot Area:

Answer Area

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

	▼
6 weeks	
12 months	
10 years	

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

	▼
6 weeks	
12 months	
10 years	

Answer Area:

Answer Area

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

	▼
6 weeks	
12 months	
10 years	

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

	▼
6 weeks	
12 months	
10 years	

Section:

Explanation:

QUESTION 9

HOTSPOT

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.

Hot Area:

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 Area:

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"/>

Section:

Explanation:

Reference:

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

QUESTION 10

HOTSPOT

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.

Hot Area:

Answer Area

RESTORE DB1 FROM

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

Answer Area:

Answer Area

RESTORE DB1 FROM

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

Section:

Explanation:

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

QUESTION 11

DRAG DROP

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

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

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

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

Select and Place:



Statements

Answer Area

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

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

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

```
CREATE CREDENTIAL
([https://mysqlbackups.blob.core.windows.net/backups])
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = 'c8a8_t0x8n0'
GO
```

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

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY =
'sqlsamplebackup'
, SECRET = 'c0yst0r4g3acc0unt4cc3ssk3y' ;
GO
```



Correct Answer:



Statements

Answer Area

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

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

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

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

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

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



Section:

Explanation:

QUESTION 12

HOTSPOT

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

You need to monitor DB 1 by using SQL Insights.

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

Hot 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**



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

Section:

Explanation:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paasoverview?view=azuresql>

QUESTION 13

You have an Azure subscription that contains an Azure SQL database named SQL1.

SQL1 is in an Azure region that does not support availability zones.

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

What should you use?

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

Correct Answer: C

Section:

QUESTION 14

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.

You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

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

Correct Answer: C

Section:

Explanation:

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



QUESTION 15

DRAG DROP

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored.

You need to ensure that SQ Server Agent jobs running on SQL 1 can notify when a failure Occurs.

Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.

Select and Place:

Actions

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



Answer Area

Correct Answer:

Actions

- Enable pager notifications upon failure.
- Create a profile named AzureManagedInstance_dbmail_profile.
-

Answer Area

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

Section:

Explanation:

QUESTION 16

DRAG DROP

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1. You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.

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

- Minimize compute times.
- Minimize storage.

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

Select and Place:

Actions

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

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

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

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

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

```
EXEC sys.sp_cdc_enable_table
```

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

```
EXEC sys.sp_cdc_enable_db
```

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

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

Answer Area



Correct Answer:



Actions

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

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

```
EXEC sys.sp_cdc_enable_table
```

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

```
EXEC sys.sp_cdc_enable_db
```

Answer Area

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

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

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

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

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

Section:

Explanation:

QUESTION 17

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`

A. Option A

B. Option B

C. Option C



D. Option D

Correct Answer: D

Section:

QUESTION 18

You have a Microsoft SQL Server 2017 server.

You need to migrate the server to Azure. The solution must meet the following requirements:

- Ensure that the latest version of SQL Server is used.
- Support the SQL Server Agent service.

Minimize administrative effort.

What should you use?

- A. SQL Server on Azure Virtual Machines
- B. Azure SQL Database
- C. an Azure SQL Database elastic pool
- D. Azure SQL Managed Instance

Correct Answer: D

Section:

QUESTION 19

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL 1 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

Correct Answer: D

Section:

QUESTION 20

HOTSPOT

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.

Hot Area:

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 Area:

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



Section:

Explanation:

QUESTION 21

HOTSPOT

You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1.

You need to configure the autogrow and autoshrink settings for DB1.

Which statements should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Autogrow:

- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options

Autoshrink:

- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options

Answer Area:

Autogrow:

- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options



Autoshrink:

- ALTER DATABASE MODIFY File and Filegroup options
- ALTER DATABASE SCOPED CONFIGURATION
- ALTER DATABASE SET options

Section:

Explanation:

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

QUESTION 22

HOTSPOT

You have an Azure SQL logical server.

You run the following script.

```

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

```

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

NOTE: Each correct selection is worth one point

Hot Area:

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

Answer Area:

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

Section:

Explanation:

QUESTION 23

HOTSPOT

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

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

CREATE USER FROM

[App1]
[Contoso\app1]
[App1@contoso.com]

login
Windows
EXTERNAL PROVIDER

Answer Area:

CREATE USER FROM

[App1]
[Contoso\app1]
[App1@contoso.com]

login
Windows
EXTERNAL PROVIDER



Section:

Explanation:

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

QUESTION 24

HOTSPOT

You need to use an Azure Resource Manager (ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O performance for the SQL Server database and log files. How should you complete the template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

```
"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,
    ...
  }
}
"resources": [
  ...
  "osDisk": {
    ...
    "copy": [
      {"name": "dataDisks","count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
      "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",
      "caching": "[if(greaterOrEquals(copyIndex('dataDisks'), parameters('dataDiskCount')),
      variables('dataDisks').caching )]", "diskSizeGB": 1023,
      ...
    ]
  }
}
```

Answer Area:



```

"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,
    ...
  }
}
"resources": [
  ...
  "osDisk": {
    ...
    "copy": [
      {
        "name": "dataDisks","count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
        "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",
        "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),parameters('dataDiskCount')),
        variables('dataDisks').caching )]", "diskSizeGB": 1023,
        
      }
    ]
  }
]

```

Section:
Explanation:

QUESTION 25

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, create a service health 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, configure an activity log alert.

Correct Answer: C

Section:

QUESTION 26

DRAG DROP

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named App1. App1 uses db1. You need to enable Resource Governor for a App1. The solution must meet the following requirements:

- App1 must be able to consume all available CPU resources.
- App1 must have at least half of the available CPU resources always available.


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

Select and Place:

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

Correct 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.

Section:

Explanation:

QUESTION 27

You have an Azure SQL database named DB1.

You need to encrypt DB1. The solution must meet the following requirements;

- Encrypt data in motion.
- Support comparison operators.
- Provide randomized encryption.

What should you include in the solution?

A. Always Encrypted

- B. column-level encryption
- C. Transparent Data Encryption (TDE)
- D. Always Encrypted with secure enclaves

Correct Answer: A

Section:

QUESTION 28

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.

Correct Answer: B

Section:

QUESTION 29

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1. You plan to migrate DB1 to an Azure SQL Database managed instance. What should you use to minimize downtime and data loss during the migration?

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



Correct Answer: D

Section:

QUESTION 30

You have an Azure subscription.

You need to deploy an Azure SQL database. The solution must meet the following requirements:

- Dynamically scale CPU resources.
- Ensure that the database can be paused to reduce costs.

What should you use?

- A. the Business Critical service tier
- B. the serverless compute tier
- C. an elastic pool
- D. the General Purpose service tier

Correct Answer: B

Section:

QUESTION 31

You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:

- Support real-time data replication to a different geographic region.
- Use Azure as a disaster recovery target.
- Minimize costs and administrative effort.

What should you include in the recommendation?

- A. database mirroring on an instance of SQL Server on Azure Virtual Machines
- B. availability groups for SQL Server on Azure Virtual Machines
- C. an Azure SQL Managed Instance link
- D. transactional replication to an Azure SQL Managed Instance

Correct Answer: D

Section:

QUESTION 32

You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1. You need to move a SQL Server Agent job from DB1 to SQLMI1. Which job attribute is unsupported in SQLMI1?

- A. log to table
- B. email notifications
- C. schedules
- D. output files

Correct Answer: D

Section:

QUESTION 33

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

You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs.

The solution must use the principle of least privilege.

Which role should you assign to User1?

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

Correct Answer: A

Section:

QUESTION 34

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



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

You plan to enable Always Encrypted for Table1.

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

NOTE: Each correct selection is worth one point

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

Correct Answer: A, D

Section:

QUESTION 35

You have a single availability set that contains two SQL Server on Azure Virtual Machines instances. The instances were deployed by using an Azure Marketplace SQL Server 2019 Enterprise image that has the latest cumulative updates applied. The instances are configured as the nodes of a failover cluster instance (FCI) named FCI1. You need to ensure that client applications can connect to FCI1. The solution must meet the following requirements:

- Provide an availability SLA
- Minimize costs.

What should you create?

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

Correct Answer: C

Section:

QUESTION 36

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

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

What should you recommend?

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

Correct Answer: A

Section:

QUESTION 37

HOTSPOT

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil. You plan to perform an online migration of db1 to MM by using Azure Database Migration Service.

You need to create the backups for the migration. The solution must minimize the number of backup files created. Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Backup type:

- Full only
- Full and differential
- Full and transaction log
- Transaction log only

To store the backups:

- Append each backup to a separate file.
- Append all the backups to a single file.
- Append each backup to a separate file.
- Write each backup to a separate file.

Answer Area:

Answer Area

Backup type:

- Full only
- Full and differential
- Full and transaction log
- Transaction log only

To store the backups:

- Append each backup to a separate file.
- Append all the backups to a single file.
- Append each backup to a separate file.
- Write each backup to a separate file.

Section:

Explanation:

QUESTION 38

HOTSPOT

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL. You need to use PowerShell to enable and configure automated patching for SQL. The solution must include both SQL Server and Windows security updates.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQL1" |
```

Set-AzVMExtension
Get-AzVMSqlServerExtension
Set-AzVMExtension
Set-AzVMSqlServerExtension

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

```
-SQLManagementType
```

-Lightweight
-Full
-Lightweight
-NoAgent

Answer Area:

Answer Area

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

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQL1" |
```

Set-AzVMExtension
Get-AzVMSqlServerExtension
Set-AzVMExtension
Set-AzVMSqlServerExtension

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

```
-SQLManagementType
```

-Lightweight
-Full
-Lightweight
-NoAgent

Section:

Explanation:

QUESTION 39

You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication. You need to migrate SQL1 to Azure. Which service should you use?

- A. Azure SQL Edge
- B. Azure SQL Database
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed instance

Correct Answer: C

Section:

QUESTION 40

You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1. You need to configure SQL! to use mixed mode authentication. Which procedure should you run?

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

Correct Answer: B

Section:

QUESTION 41

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory. You scale up the virtual machine to 8 vCPUs and 64 GB of memory. You need to reduce tempdb contention without negatively affecting server performance. What is the number of secondary data files that you should configure for tempdb?

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

Correct Answer: C

Section:

QUESTION 42

DRAG DROP

You have two instances of SQL Server on Azure virtual Machines named VM1 and VM2. VM1 hosts a database named db1. You plan to create a database availability group (DAG) for db1. The solution must use certificate authentication between VM1 and VM2. You need to configure authentication for the outbound connections of 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.

Select and Place:

Actions

- From the master database, run CREATE_ENDPOINT.
- From db1, run CREATE_CERTIFICATE.
- From the master database, run CREATE_MASTER_KEY_ENCRYPTION_BY_PASSWORD.
- From the master database, run CREATE_CERTIFICATE.
- From db1, run CREATE_ENDPOINT.

Answer Area

-
-
-

Correct Answer:

Actions

- From the master database, run CREATE_ENDPOINT.
- From db1, run CREATE_CERTIFICATE.
-
-

Answer Area

- From the master database, run CREATE_MASTER_KEY_ENCRYPTION_BY_PASSWORD.
- From the master database, run CREATE_CERTIFICATE.
- From db1, run CREATE_ENDPOINT.

Section:
Explanation:

QUESTION 43
HOTSPOT

You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1. You need to identify and repair any physical or logical corruption in DB1. The solution must meet the following requirements:

- Minimize how long it takes to complete the procedure.
- Minimize data loss.

How should you complete the command? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

DBCC CHECK [DB1],

REPAIR_REBUILD) WITH

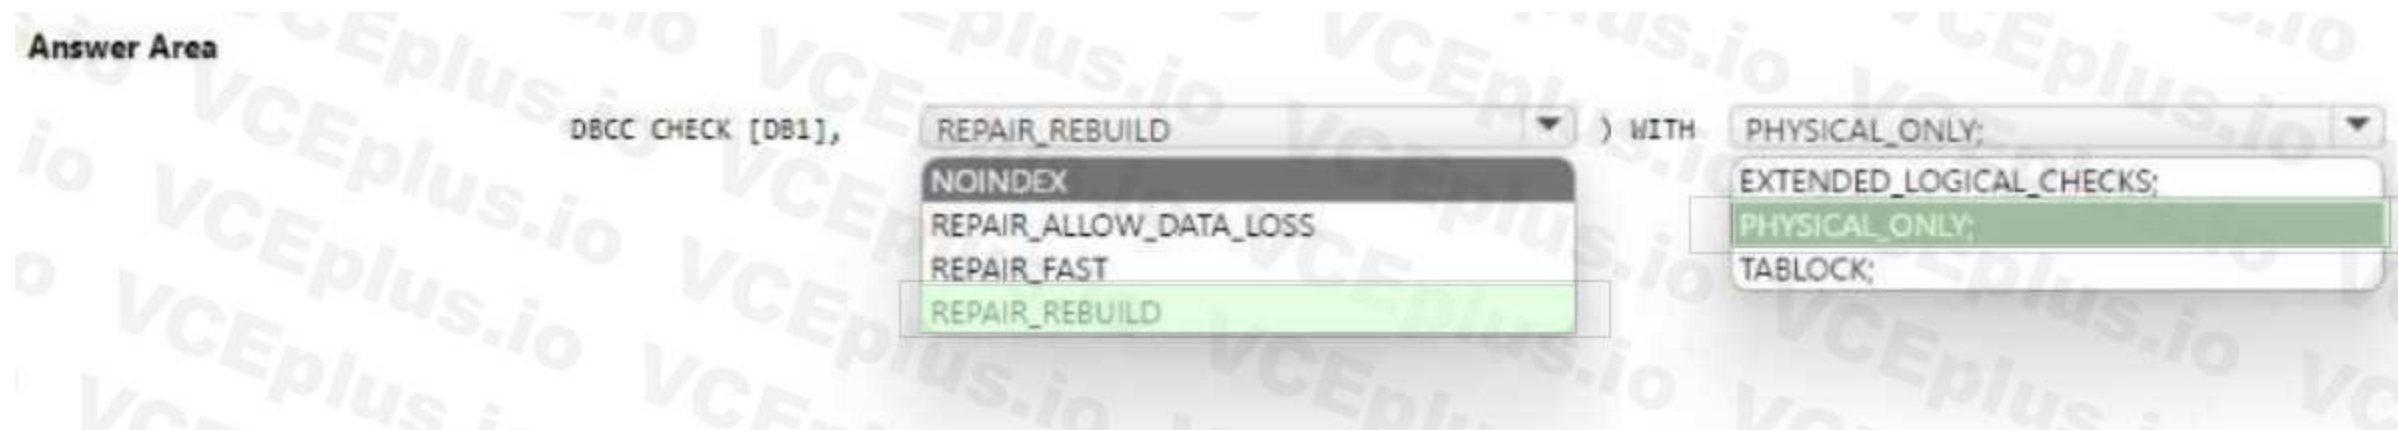
PHYSICAL_ONLY;

EXTENDED_LOGICAL_CHECKS;

PHYSICAL_ONLY;

TABLOCK;

Answer Area:



Section:

Explanation:

QUESTION 44

You have an Azure subscription.

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

- Minimize how long it takes to fail over.
- Maintain existing connections to the primary replica during a failover.

What should you do?

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

Correct Answer: C

Section:

QUESTION 45

HOTSPOT

You have an Azure SQL database named DB1 that contains a table named Table 1.

You run a query to load data into Table1.

The performance metrics of Table1 during the load operation are 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.

Hot Area:

Answer Area

To reduce how long it takes to complete the query you must [answer choice].

use an elastic pool
scale the resource
use an elastic pool
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

an In-Memory OTLP non durable
a temporary
an In-Memory OTLP durable
an In-Memory OTLP non durable

Answer Area:

Answer Area

To reduce how long it takes to complete the query you must [answer choice].

use an elastic pool
scale the resource
use an elastic pool
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

an In-Memory OTLP non durable
a temporary
an In-Memory OTLP durable
an In-Memory OTLP non durable



Section:

Explanation:

QUESTION 46

HOTSPOT

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. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To collect monitoring data, use:

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

To store monitoring data, create:

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

Answer Area:

Answer Area

To collect monitoring data, use:

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

To store monitoring data, create:

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

Section:

Explanation:

QUESTION 47

HOTSPOT

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

The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:

- Capture raw event data and store the data in Azure Storage.
- Minimize the performance impact of capturing extended events.

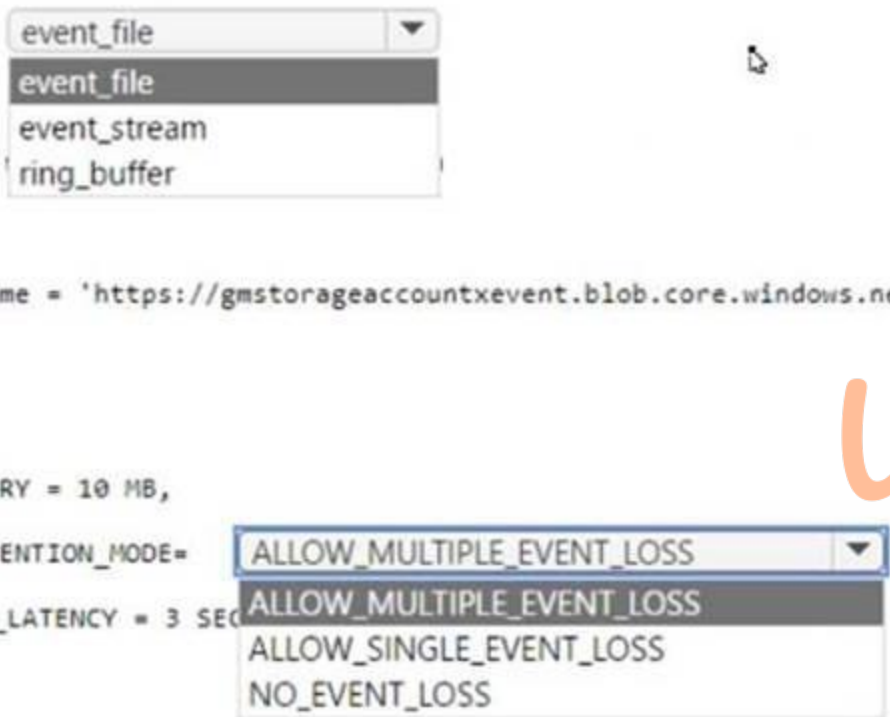
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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



Answer Area:

Answer Area

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



The screenshot shows the configuration of an event session in SQL Server Enterprise Manager. The 'package0.' dropdown is set to 'event_file'. The 'EVENT_RETENTION_MODE=' dropdown is set to 'ALLOW_MULTIPLE_EVENT_LOSS'. The 'MAX_DISPATCH_LATENCY = 3 SEC' is also visible.

Section:

Explanation:

QUESTION 48

You have an Azure SQL database.

You need to implement a disaster recovery solution that meets the following requirements:

- Minimizes how long it takes to recover the database if a datacenter fails
- Minimizes administrative effort

What should you include in the solution?

- A. Azure Backup
- B. active geo-replication
- C. Azure Site Recovery
- D. auto-failover groups

Correct Answer: C

Section:

QUESTION 49

HOTSPOT

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.

Hot Area:

Answer Area

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



Answer Area:

Answer Area

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

Section:

Explanation:

QUESTION 50

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

Name	Type	Description
SQLSVR1	Logical SQL server	Hosts one Azure SQL database
SQLDB11	Azure SQL database	Is hosted by SQLSVR1
SQLSVR2	Logical SQL server	Hosts two Azure SQL databases
SQLDB21	Azure SQL database	Is hosted by SQLSVR2
SQLDB22	Azure SQL database	Is hosted by SQLSVR2

You plan to use SQLDB11 as an elastic job database to run jobs on SQLDB11 and SQLDB22. What is the minimum number of database scoped credentials required for the elastic jobs?

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

Correct Answer: A

Section:

