

**DP-200.72q**

Number: DP-200  
Passing Score: 800  
Time Limit: 120 min

**DP-200**



**Website:** <https://vceplus.com>

**VCE to PDF Converter:** <https://vceplus.com/vce-to-pdf/>

**Facebook:** <https://www.facebook.com/VCE.For.All.VN/>

**Twitter :** [https://twitter.com/VCE\\_Plus](https://twitter.com/VCE_Plus)

<https://vceplus.com/>

### **Implementing an Azure Data Solution**

#### **Question Set 1**

#### **QUESTION 1**

You develop data engineering solutions for a company. The company has on-premises Microsoft SQL Server databases at multiple locations.

<https://vceplus.com/>

The company must integrate data with Microsoft Power BI and Microsoft Azure Logic Apps. The solution must avoid single points of failure during connection and transfer to the cloud. The solution must also minimize latency.

You need to secure the transfer of data between on-premises databases and Microsoft Azure.



<https://vceplus.com/> What should

you do?

- A. Install a standalone on-premises Azure data gateway at each location
- B. Install an on-premises data gateway in personal mode at each location
- C. Install an Azure on-premises data gateway at the primary location
- D. Install an Azure on-premises data gateway as a cluster at each location

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You can create high availability clusters of On-premises data gateway installations, to ensure your organization can access on-premises data resources used in Power BI reports and dashboards. Such clusters allow gateway administrators to group gateways to avoid single points of failure in accessing on-premises data resources. The Power BI service always uses the primary gateway in the cluster, unless it's not available. In that case, the service switches to the next gateway in the cluster, and so on.

References: <https://docs.microsoft.com/en-us/power-bi/service-gateway-high-availability-clusters>

## QUESTION 2

You are a data architect. The data engineering team needs to configure a synchronization of data between an on-premises Microsoft SQL Server database to Azure SQL Database.

Ad-hoc and reporting queries are being overutilized the on-premises production instance. The synchronization process must:

<https://vceplus.com/>

- Perform an initial data synchronization to Azure SQL Database with minimal downtime ▪
- Perform bi-directional data synchronization after initial synchronization

You need to implement this synchronization solution.

Which synchronization method should you use?

- A. transactional replication
- B. Data Migration Assistant (DMA)
- C. backup and restore
- D. SQL Server Agent job
- E. Azure SQL Data Sync

**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

SQL Data Sync is a service built on Azure SQL Database that lets you synchronize the data you select bi-directionally across multiple SQL databases and SQL Server instances.

With Data Sync, you can keep data synchronized between your on-premises databases and Azure SQL databases to enable hybrid applications.

Compare Data Sync with Transactional Replication

	<b>Data Sync</b>	<b>Transactional Replication</b>
Advantages	<ul style="list-style-type: none"><li>- Active-active support</li><li>- Bi-directional between on-premises and Azure SQL Database</li></ul>	<ul style="list-style-type: none"><li>- Lower latency</li><li>- Transactional consistency</li><li>- Reuse existing topology after migration</li></ul>
Disadvantages	<ul style="list-style-type: none"><li>- 5 min or more latency</li><li>- No transactional consistency</li><li>- Higher performance impact</li></ul>	<ul style="list-style-type: none"><li>- Can't publish from Azure SQL Database single database or pooled database</li><li>- High maintenance cost</li></ul>

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

### QUESTION 3

An application will use Microsoft Azure Cosmos DB as its data solution. The application will use the Cassandra API to support a column-based database type that uses containers to store items.

You need to provision Azure Cosmos DB. Which container name and item name should you use? Each correct answer presents part of the solutions.

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

- A. collection
- B. rows
- C. graph
- D. entities
- E. table

**Correct Answer:** BE

**Section:** [none]

**Explanation**



**Explanation/Reference:**

Explanation:

B: Depending on the choice of the API, an Azure Cosmos item can represent either a document in a collection, a row in a table or a node/edge in a graph. The following table shows the mapping between API-specific entities to an Azure Cosmos item:

Cosmos entity	SQL API	Cassandra API	Azure Cosmos DB's API for MongoDB	Gremlin API	Table API
Azure Cosmos item	Document	Row	Document	Node or Edge	Item

E: An Azure Cosmos container is specialized into API-specific entities as follows:

Azure Cosmos entity	SQL API	Cassandra API	Azure Cosmos DB's API for MongoDB	Gremlin API	Table API
Azure Cosmos container	Collection	Table	Collection	Graph	Table

References: <https://docs.microsoft.com/en-us/azure/cosmos-db/databases-containers-items>

#### QUESTION 4

A company has a SaaS solution that uses Azure SQL Database with elastic pools. The solution contains a dedicated database for each customer organization. Customer organizations have peak usage at different periods during the year.

You need to implement the Azure SQL Database elastic pool to minimize cost.

Which option or options should you configure?

- A. Number of transactions only
- B. eDTUs per database only
- C. Number of databases only
- D. CPU usage only
- E. eDTUs and max data size



**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

The best size for a pool depends on the aggregate resources needed for all databases in the pool. This involves determining the following:

- Maximum resources utilized by all databases in the pool (either maximum DTUs or maximum vCores depending on your choice of resourcing model).
- Maximum storage bytes utilized by all databases in the pool.

Note: Elastic pools enable the developer to purchase resources for a pool shared by multiple databases to accommodate unpredictable periods of usage by individual databases. You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.

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

### QUESTION 5

The data engineering team manages Azure HDInsight clusters. The team spends a large amount of time creating and destroying clusters daily because most of the data pipeline process runs in minutes.

You need to implement a solution that deploys multiple HDInsight clusters with minimal effort.

What should you implement?

- A. Azure Databricks
- B. Azure Traffic Manager
- C. Azure Resource Manager templates
- D. Ambari web user interface

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

A Resource Manager template makes it easy to create the following resources for your application in a single, coordinated operation: ▪ HDInsight clusters and their dependent resources (such as the default storage account). ▪ Other resources (such as Azure SQL Database to use Apache Sqoop).

In the template, you define the resources that are needed for the application. You also specify deployment parameters to input values for different environments. The template consists of JSON and expressions that you use to construct values for your deployment.

References: <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-create-linux-clusters-arm-templates>

### QUESTION 6

You are the data engineer for your company. An application uses a NoSQL database to store data. The database uses the key-value and wide-column NoSQL database type.

Developers need to access data in the database using an API.

You need to determine which API to use for the database model and type.

Which two APIs should you use? Each correct answer presents a complete solution.

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

- A. Table API
- B. MongoDB API
- C. Gremlin API
- D. SQL API
- E. Cassandra API

**Correct Answer:** BE

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

B: Azure Cosmos DB is the globally distributed, multimodel database service from Microsoft for mission-critical applications. It is a multimodel database and supports document, key-value, graph, and columnar data models.

E: Wide-column stores store data together as columns instead of rows and are optimized for queries over large datasets. The most popular are Cassandra and HBase.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction>

<https://www.mongodb.com/scale/types-of-nosql-databases>



### QUESTION 7

A company is designing a hybrid solution to synchronize data and on-premises Microsoft SQL Server database to Azure SQL Database.

You must perform an assessment of databases to determine whether data will move without compatibility issues. You need to perform the assessment.

Which tool should you use?

- A. SQL Server Migration Assistant (SSMA)
- B. Microsoft Assessment and Planning Toolkit
- C. SQL Vulnerability Assessment (VA)
- D. Azure SQL Data Sync
- E. Data Migration Assistant (DMA)

**Correct Answer:** E

**Section:** [none]

**Explanation**

<https://vceplus.com/>

**Explanation/Reference:**

Explanation:

The Data Migration Assistant (DMA) helps you upgrade to a modern data platform by detecting compatibility issues that can impact database functionality in your new version of SQL Server or Azure SQL Database. DMA recommends performance and reliability improvements for your target environment and allows you to move your schema, data, and uncontained objects from your source server to your target server.

References: <https://docs.microsoft.com/en-us/sql/dma/dma-overview>

**QUESTION 8**

A company plans to use Azure SQL Database to support a mission-critical application.

The application must be highly available without performance degradation during maintenance windows.

You need to implement the solution.

Which three technologies should you implement? Each correct answer presents part of the solution.

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

- A. Premium service tier
- B. Virtual machine Scale Sets
- C. Basic service tier
- D. SQL Data Sync
- E. Always On availability groups
- F. Zone-redundant configuration



**Correct Answer:** AEF

**Section:** [none]

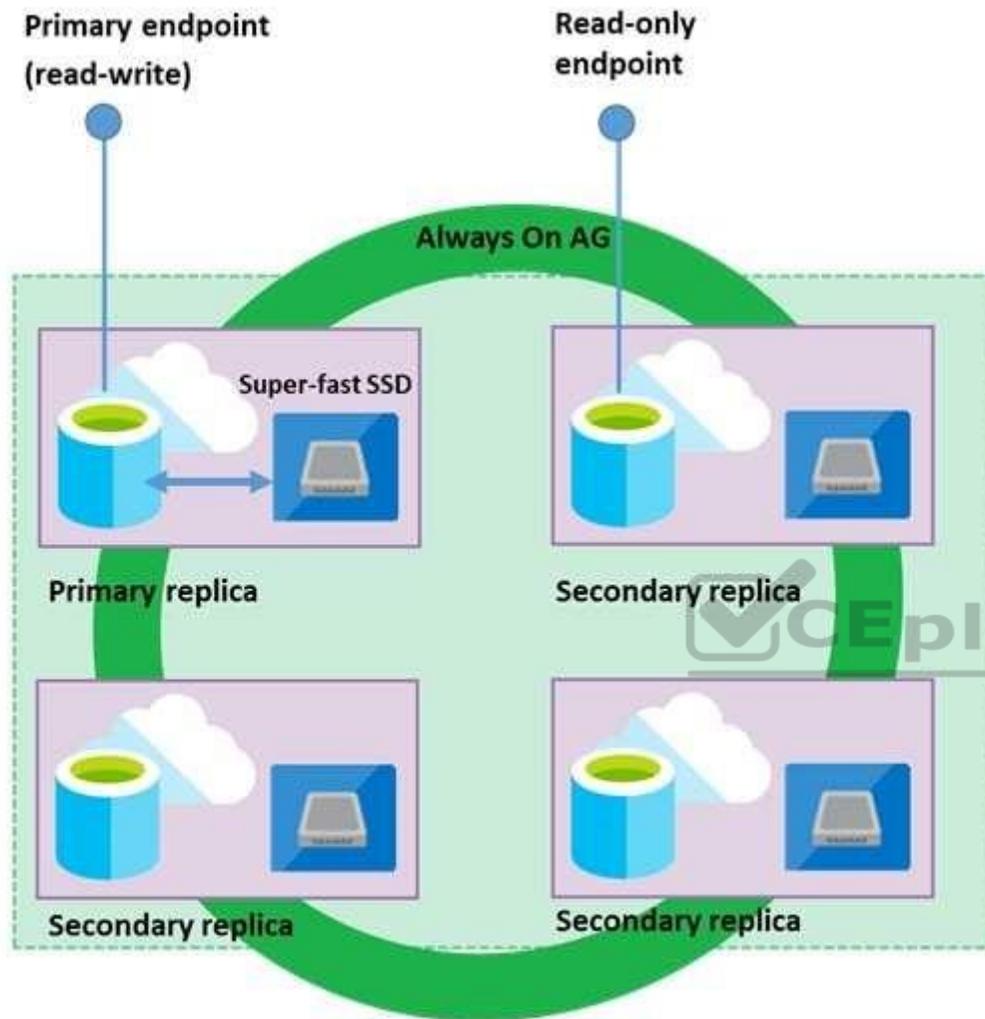
**Explanation**

**Explanation/Reference:**

Explanation:

A: Premium/business critical service tier model that is based on a cluster of database engine processes. This architectural model relies on a fact that there is always a quorum of available database engine nodes and has minimal performance impact on your workload even during maintenance activities.

E: In the premium model, Azure SQL database integrates compute and storage on the single node. High availability in this architectural model is achieved by replication of compute (SQL Server Database Engine process) and storage (locally attached SSD) deployed in 4-node cluster, using technology similar to SQL Server Always On Availability Groups.



Business Critical service tier: collocated compute and storage

F: Zone redundant configuration By default, the quorum-set replicas for the local storage configurations are created in the same datacenter. With the introduction of Azure Availability Zones, you have the ability to place the different replicas in the quorum-sets to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

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

### QUESTION 9

A company plans to use Azure Storage for file storage purposes. Compliance rules require:

- A single storage account to store all operations including reads, writes and deletes
- Retention of an on-premises copy of historical operations

You need to configure the storage account.

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

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

- A. Configure the storage account to log read, write and delete operations for service type Blob
- B. Use the AzCopy tool to download log data from \$logs/blob
- C. Configure the storage account to log read, write and delete operations for service-type table
- D. Use the storage client to download log data from \$logs/table
- E. Configure the storage account to log read, write and delete operations for service type queue

**Correct Answer:** AB

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Storage Logging logs request data in a set of blobs in a blob container named \$logs in your storage account. This container does not show up if you list all the blob containers in your account but you can see its contents if you access it directly.

To view and analyze your log data, you should download the blobs that contain the log data you are interested in to a local machine. Many storage-browsing tools enable you to download blobs from your storage account; you can also use the Azure Storage team provided command-line Azure Copy Tool (AzCopy) to download your log data.

References:

<https://docs.microsoft.com/en-us/rest/api/storageservices/enabling-storage-logging-and-accessing-log-data>

**QUESTION 10**  
DRAG DROP

You are developing a solution to visualize multiple terabytes of geospatial data.

The solution has the following requirements:

- Data must be encrypted.
- Data must be accessible by multiple resources on Microsoft Azure.

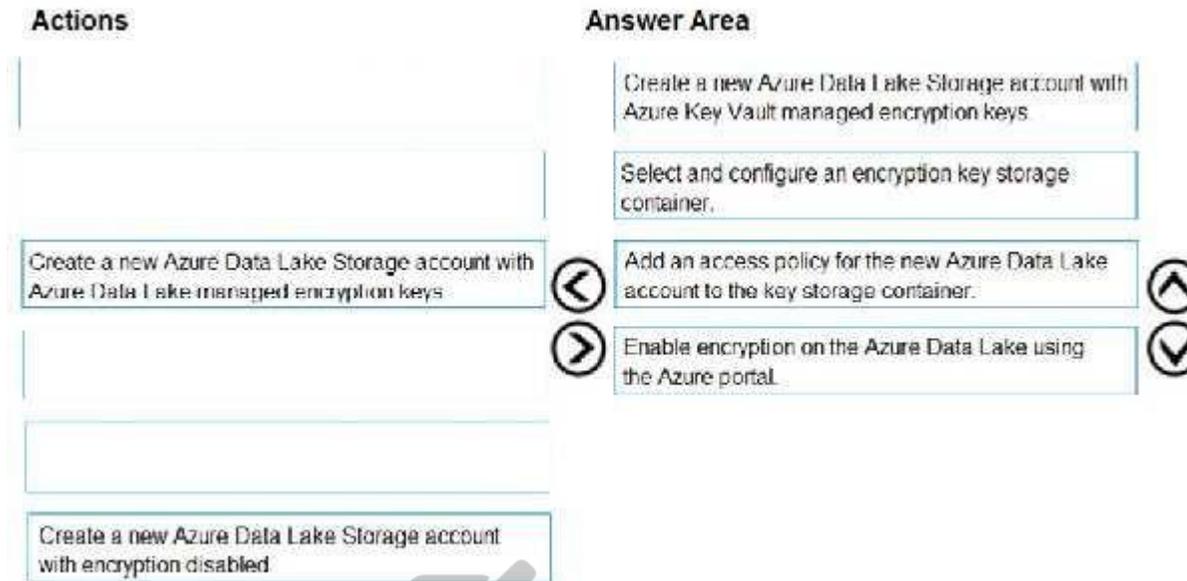
You need to provision storage for the solution.

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

**Select and Place:**

Actions	Answer Area
Enable encryption on the Azure Data Lake using the Azure portal.	
Add an access policy for the new Azure Data Lake account to the key storage container.	
Create a new Azure Data Lake Storage account with Azure Data Lake managed encryption keys.	⏪
Select and configure an encryption key storage container.	⏩
Create a new Azure Data Lake Storage account with Azure Key Vault managed encryption keys.	⏴
Create a new Azure Data Lake Storage account with encryption disabled.	⏵

**Correct Answer:**



**Section:** [none]  
**Explanation**

**Explanation/Reference:**  
 Explanation:

Create a new Azure Data Lake Storage account with Azure Data Lake managed encryption keys

For Azure services, Azure Key Vault is the recommended key storage solution and provides a common management experience across services. Keys are stored and managed in key vaults, and access to a key vault can be given to users or services. Azure Key Vault supports customer creation of keys or import of customer keys for use in customer-managed encryption key scenarios.

Note: Data Lake Storage Gen1 account Encryption Settings. There are three options:

- Do not enable encryption.
- Use keys managed by Data Lake Storage Gen1, if you want Data Lake Storage Gen1 to manage your encryption keys.
- Use keys from your own Key Vault. You can select an existing Azure Key Vault or create a new Key Vault. To use the keys from a Key Vault, you must assign permissions for the Data Lake Storage Gen1 account to access the Azure Key Vault.

References: <https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-atrest> **QUESTION 11**

You are developing a data engineering solution for a company. The solution will store a large set of key-value pair data by using Microsoft Azure Cosmos DB.

<https://vceplus.com/>

The solution has the following requirements:

- Data must be partitioned into multiple containers.
  - Data containers must be configured separately.
  - Data must be accessible from applications hosted around the world. ▪
- The solution must minimize latency.

You need to provision Azure Cosmos DB.

- A. Cosmos account-level throughput.
- B. Provision an Azure Cosmos DB account with the Azure Table API. Enable geo-redundancy.
- C. Configure table-level throughput.
- D. Replicate the data globally by manually adding regions to the Azure Cosmos DB account.
- E. Provision an Azure Cosmos DB account with the Azure Table API. Enable multi-region writes.

**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Scale read and write throughput globally. You can enable every region to be writable and elastically scale reads and writes all around the world. The throughput that your application configures on an Azure Cosmos database or a container is guaranteed to be delivered across all regions associated with your Azure Cosmos account. The provisioned throughput is guaranteed up by financially backed SLAs.

References: <https://docs.microsoft.com/en-us/azure/cosmos-db/distribute-data-globally>

## QUESTION 12

A company has a SaaS solution that uses Azure SQL Database with elastic pools. The solution will have a dedicated database for each customer organization. Customer organizations have peak usage at different periods during the year.

Which two factors affect your costs when sizing the Azure SQL Database elastic pools? Each correct answer presents a complete solution.

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

- A. maximum data size
- B. number of databases
- C. eDTUs consumption
- D. number of read operations

<https://vceplus.com/>

E. number of transactions

**Correct Answer:** AC

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

A: With the vCore purchase model, in the General Purpose tier, you are charged for Premium blob storage that you provision for your database or elastic pool. Storage can be configured between 5 GB and 4 TB with 1 GB increments. Storage is priced at GB/month.

C: In the DTU purchase model, elastic pools are available in basic, standard and premium service tiers. Each tier is distinguished primarily by its overall performance, which is measured in elastic Database Transaction Units (eDTUs).

References:

<https://azure.microsoft.com/en-in/pricing/details/sql-database/elastic/>

### QUESTION 13

A company runs Microsoft SQL Server in an on-premises virtual machine (VM).

You must migrate the database to Azure SQL Database. You synchronize users from Active Directory to Azure Active Directory (Azure AD).

You need to configure Azure SQL Database to use an Azure AD user as administrator.

What should you configure?

- A. For each Azure SQL Database, set the Access Control to administrator.
- B. For each Azure SQL Database server, set the Active Directory to administrator.
- C. For each Azure SQL Database, set the Active Directory administrator role.
- D. For each Azure SQL Database server, set the Access Control to administrator.

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

There are two administrative accounts (Server admin and Active Directory admin) that act as administrators.

One Azure Active Directory account, either an individual or security group account, can also be configured as an administrator. It is optional to configure an Azure AD administrator, but an Azure AD administrator must be configured if you want to use Azure AD accounts to connect to SQL Database.

<https://vceplus.com/>

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

#### QUESTION 14

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have an Azure SQL database named DB1 that contains a table named Table1. Table1 has a field named Customer\_ID that is varchar(22).

You need to implement masking for the Customer\_ID field to meet the following requirements:

- The first two prefix characters must be exposed.
  - The last four prefix characters must be exposed. ▪
- All other characters must be masked.

Solution: You implement data masking and use an email function mask.

Does this meet the goal?

- A. Yes
- B. No



**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Must use Custom Text data masking, which exposes the first and last characters and adds a custom padding string in the middle.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started>

#### QUESTION 15

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have an Azure SQL database named DB1 that contains a table named Table1. Table1 has a field named Customer\_ID that is varchar(22).

<https://vceplus.com/>

You need to implement masking for the Customer\_ID field to meet the following requirements:

- The first two prefix characters must be exposed.
  - The last four prefix characters must be exposed. ▪
- All other characters must be masked.

Solution: You implement data masking and use a random number function mask.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Must use Custom Text data masking, which exposes the first and last characters and adds a custom padding string in the middle.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started>

#### **QUESTION 16**

You plan to create a dimension table in Azure Data Warehouse that will be less than 1 GB.

You need to create the table to meet the following requirements:

- Provide the fastest query time. ▪
- Minimize data movement.

Which type of table should you use?

- A. hash distributed
- B. heap
- C. replicated
- D. round-robin

**Correct Answer:** D

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Usually common dimension tables or tables that doesn't distribute evenly are good candidates for round-robin distributed table.

Note: Dimension tables or other lookup tables in a schema can usually be stored as round-robin tables. Usually these tables connect to more than one fact tables and optimizing for one join may not be the best idea. Also usually dimension tables are smaller which can leave some distributions empty when hash distributed. Round-robin by definition guarantees a uniform data distribution.

References: <https://blogs.msdn.microsoft.com/sqlcat/2015/08/11/choosing-hash-distributed-table-vs-round-robin-distributed-table-in-azure-sql-dw-service/>

**QUESTION 17**

You have an Azure SQL data warehouse.

Using PolyBase, you create table named [Ext].[Items] to query Parquet files stored in Azure Data Lake Storage Gen2 without importing the data to the data warehouse.

The external table has three columns.



You discover that the Parquet files have a fourth column named ItemID.

Which command should you run to add the ItemID column to the external table?

A.

```
DROP TABLE [Ext].[Items]
CREATE EXTERNAL TABLE [Ext].[Items]_
( [ItemID] [int] NULL,
  [ItemName] nvarchar(50) NULL,
  [ItemType] nvarchar(20) NULL,
  [ItemDescription] nvarchar(250))
WITH
(
  LOCATION='/Items/',
  DATA_SOURCE = AzureDataLakeStore,
  FILE_FORMAT = PARQUET,
  REJECT_TYPE = VALUE,
  REJECT_VALUE = 0
);
```

B.

```
ALTER EXTERNAL TABLE [Ext].[Items]
ADD [ItemID] int;
```



C.

```
DROP EXTERNAL FILE FORMAT parquetfile1;
CREATE EXTERNAL FILE FORMAT parquetfile1
WITH (
  FORMAT_TYPE = PARQUET,
  DATA_COMPRESSION = 'org.apache.hadoop.io.compress.SnappyCodec'
);
```

D.

```
ALTER TABLE [Ext].[Items]
ADD [ItemID] int
```

- A. Option A
- B. Option B
- C. Option C

D. Option D

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Incorrect Answers:

B, D: Only these Data Definition Language (DDL) statements are allowed on external tables:

- CREATE TABLE and DROP TABLE
- CREATE STATISTICS and DROP STATISTICS ▪
- CREATE VIEW and DROP VIEW

References:

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

#### **QUESTION 18**

**DRAG DROP**

You have a table named SalesFact in an Azure SQL data warehouse. SalesFact contains sales data from the past 36 months and has the following characteristics:

- Is partitioned by month
- Contains one billion rows
- Has clustered columnstore indexes

All the beginning of each month, you need to remove data SalesFact that is older than 36 months as quickly as possible.

Which three actions should you perform in sequence in a stored procedure? 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 an empty table named SalesFact_Work that has the same schema as SalesFact.	
Drop the SalesFact_Work table.	
Copy the data to a new table by using CREATE TABLE AS SELECT (CTAS).	
Truncate the partition containing the stale data.	
Switch the partition containing the stale data from SalesFact to SalesFact_Work.	
Execute a DELETE statement where the value in the Date column is more than 36 months ago.	

**Correct Answer:**

Actions	Answer Area
	Create an empty table named SalesFact_Work that has the same schema as SalesFact.
	Switch the partition containing the stale data from SalesFact to SalesFact_Work.
Copy the data to a new table by using CREATE TABLE AS SELECT (CTAS).	Drop the SalesFact_Work table.
Truncate the partition containing the stale data.	
Execute a DELETE statement where the value in the Date column is more than 36 months ago.	

**Section:** [none]  
**Explanation**

**Explanation/Reference:**

Step 1: Create an empty table named SalesFact\_work that has the same schema as SalesFact.

<https://vceplus.com/>

Step 2: Switch the partition containing the stale data from SalesFact to SalesFact\_Work.

SQL Data Warehouse supports partition splitting, merging, and switching. To switch partitions between two tables, you must ensure that the partitions align on their respective boundaries and that the table definitions match.

Loading data into partitions with partition switching is a convenient way stage new data in a table that is not visible to users the switch in the new data.

Step 3: Drop the SalesFact\_Work table.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-partition>

### QUESTION 19

You plan to implement an Azure Cosmos DB database that will write 100,000 JSON every 24 hours. The database will be replicated to three regions. Only one region will be writable.

You need to select a consistency level for the database to meet the following requirements:

- Guarantee monotonic reads and writes within a session.
- Provide the fastest throughput. ▪ Provide the lowest latency.

Which consistency level should you select?

- A. Strong
- B. Bounded Staleness
- C. Eventual
- D. Session
- E. Consistent Prefix

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Session: Within a single client session reads are guaranteed to honor the consistent-prefix (assuming a single “writer” session), monotonic reads, monotonic writes, read-your-writes, and write-follows-reads guarantees. Clients outside of the session performing writes will see eventual consistency.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>



## QUESTION 20

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have an Azure SQL database named DB1 that contains a table named Table1. Table1 has a field named Customer\_ID that is varchar(22).

You need to implement masking for the Customer\_ID field to meet the following requirements:

- The first two prefix characters must be exposed.
  - The last four prefix characters must be exposed. ▪
- All other characters must be masked.

Solution: You implement data masking and use a credit card function mask.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

We must use Custom Text data masking, which exposes the first and last characters and adds a custom padding string in the middle.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started>

## QUESTION 21

**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 scenario, 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 Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB.

You plan to copy the data from the storage account to an Azure SQL data warehouse.

<https://vceplus.com/>

You need to prepare the files to ensure that the data copies quickly.

Solution: You modify the files to ensure that each row is more than 1 MB.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B  
**Section:** [none]  
**Explanation**

**Explanation/Reference:**

Explanation:

Instead modify the files to ensure that each row is less than 1 MB.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>

## QUESTION 22

**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 scenario, 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 Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB.

You plan to copy the data from the storage account to an Azure SQL data warehouse.

You need to prepare the files to ensure that the data copies quickly.

Solution: You copy the files to a table that has a columnstore index.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

<https://vceplus.com/>

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Instead modify the files to ensure that each row is less than 1 MB.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>



## Testlet 2

### Case Study

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

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

Litware, Inc, is an international car racing and manufacturing company that has 1,000 employees. Most employees are located in Europe. The company supports racing teams that complete in a worldwide racing series.

#### Physical Locations

Litware has two main locations: a main office in London, England, and a manufacturing plant in Berlin, Germany.

During each race weekend, 100 engineers set up a remote portable office by using a VPN to connect the datacentre in the London office. The portable office is set up and torn down in approximately 20 different countries each year.

#### Existing environment

##### Race Central

During race weekends, Litware uses a primary application named Race Central. Each car has several sensors that send real-time telemetry data to the London datacentre. The data is used for real-time tracking of the cars.

Race Central also sends batch updates to an application named Mechanical Workflow by using Microsoft SQL Server Integration Services (SSIS).

<https://vceplus.com/>

The telemetry data is sent to a MongoDB database. A custom application then moves the data to databases in SQL Server 2017. The telemetry data in MongoDB has more than 500 attributes. The application changes the attribute names when the data is moved to SQL Server 2017.

The database structure contains both OLAP and OLTP databases.

### **Mechanical Workflow**

Mechanical Workflow is used to track changes and improvements made to the cars during their lifetime.

Currently, Mechanical Workflow runs on SQL Server 2017 as an OLAP system.

Mechanical Workflow has a named Table1 that is 1 TB. Large aggregations are performed on a single column of Table 1.

### **Requirements**

#### **Planned Changes**

Litware is the process of rearchitecting its data estate to be hosted in Azure. The company plans to decommission the London datacentre and move all its applications to an Azure datacentre.

#### **Technical Requirements**

Litware identifies the following technical requirements:

- Data collection for Race Central must be moved to Azure Cosmos DB and Azure SQL Database. The data must be written to the Azure datacentre closest to each race and must converge in the least amount of time.
- The query performance of Race Central must be stable, and the administrative time it takes to perform optimizations must be minimized.
- The datacentre for Mechanical Workflow must be moved to Azure SQL data Warehouse.
- Transparent data encryption (TDE) must be enabled on all data stores, whenever possible.
- An Azure Data Factory pipeline must be used to move data from Cosmos DB to SQL Database for Race Central. If the data load takes longer than 20 minutes, configuration changes must be made to Data Factory.
- The telemetry data must migrate toward a solution that is native to Azure.
- The telemetry data must be monitored for performance issues. You must adjust the Cosmos DB Request Units per second (RU/s) to maintain a performance SLA while minimizing the cost of the Ru/s.

#### **Data Masking Requirements**

During rare weekends, visitors will be able to enter the remote portable offices. Litware is concerned that some proprietary information might be exposed. The company identifies the following data masking requirements for the Race Central data that will be stored in SQL Database:

- Only show the last four digits of the values in a column named SuspensionSprings.

- Only Show a zero value for the values in a column named ShockOilWeight.

### QUESTION 1

On which data store you configure TDE to meet the technical requirements?

- A. Cosmos DB
- B. SQL Data Warehouse
- C. SQL Database

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Scenario: Transparent data encryption (TDE) must be enabled on all data stores, whenever possible. The datacentre for Mechanical Workflow must be moved to Azure SQL data Warehouse.

Incorrect Answers:

A: Cosmos DB does not support TDE.



### QUESTION 2

HOTSPOT

You are building the data store solution for Mechanical Workflow.

How should you configure Table1? To answer, select the appropriate options in the answer area.

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

**Hot Area:**

Answer Area

Table Type:	<input type="text"/>
	▼
	Hash distributed
	Replicated
	Round-robin

Index type:	<input type="text"/>
	▼
	Clustered
	Clustered columnstore
	Heap
	Nonclustered

Correct Answer:

Answer Area

Table Type:	<input type="text"/>
	▼
	Hash distributed
	Replicated
	Round-robin

Index type:	<input type="text"/>
	▼
	Clustered
	Clustered columnstore
	Heap
	Nonclustered

Section: [none]

Explanation

**Explanation/Reference:**

Table Type: Hash distributed.

Hash-distributed tables improve query performance on large fact tables.

Index type: Clustered columnstore

Scenario:

<https://vceplus.com/>

Mechanical Workflow has a named Table1 that is 1 TB. Large aggregations are performed on a single column of Table 1.

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



## Question Set 1

### QUESTION 1

A company uses Azure SQL Database to store sales transaction data. Field sales employees need an offline copy of the database that includes last year's sales on their laptops when there is no internet connection available.

You need to create the offline export copy.

Which three options can you use? Each correct answer presents a complete solution.

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

- A. Export to a BACPAC file by using Azure Cloud Shell, and save the file to an Azure storage account
- B. Export to a BACPAC file by using SQL Server Management Studio. Save the file to an Azure storage account
- C. Export to a BACPAC file by using the Azure portal
- D. Export to a BACPAC file by using Azure PowerShell and save the file locally
- E. Export to a BACPAC file by using the SqlPackage utility

**Correct Answer:** BCE

**Section:** [none]

**Explanation**



#### **Explanation/Reference:**

Explanation:

You can export to a BACPAC file using the Azure portal.

You can export to a BACPAC file using SQL Server Management Studio (SSMS). The newest versions of SQL Server Management Studio provide a wizard to export an Azure SQL database to a BACPAC file.

You can export to a BACPAC file using the SQLPackage utility.

Incorrect Answers:

D: You can export to a BACPAC file using PowerShell. Use the New-AzSqlDatabaseExport cmdlet to submit an export database request to the Azure SQL Database service. Depending on the size of your database, the export operation may take some time to complete. However, the file is not stored locally.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-export>

### QUESTION 2

**Note:** This question is part of series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Create a remote service binding pointing to the Azure Data Lake Gen 2 storage account
2. Create an external file format and external table using the external data source
3. Load the data using the `CREATE TABLE AS SELECT` statement Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You need to create an external file format and external table from an external data source, instead from a remote service binding pointing.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

### QUESTION 3

**Note: This question is part of series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.**

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Create an external data source pointing to the Azure storage account
2. Create a workload group using the Azure storage account name as the pool name
3. Load the data using the `CREATE TABLE AS SELECT` statement

Does the solution meet the goal?

- A. Yes

B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Use the Azure Data Lake Gen 2 storage account.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

#### QUESTION 4

You need to develop a pipeline for processing data. The pipeline must meet the following requirements:

- Scale up and down resources for cost reduction
- Use an in-memory data processing engine to speed up ETL and machine learning operations.
- Use streaming capabilities
- Provide the ability to code in SQL, Python, Scala, and R
- Integrate workspace collaboration with Git What should you use?



- A. HDInsight Spark Cluster
- B. Azure Stream Analytics
- C. HDInsight Hadoop Cluster
- D. Azure SQL Data Warehouse
- E. HDInsight Kafka Cluster
- F. HDInsight Storm Cluster

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Apache Spark is an open-source, parallel-processing framework that supports in-memory processing to boost the performance of big-data analysis applications.

<https://vceplus.com/>

HDInsight is a managed Hadoop service. Use it to deploy and manage Hadoop clusters in Azure. For batch processing, you can use Spark, Hive, Hive LLAP, MapReduce.

Languages: R, Python, Java, Scala, SQL

You can create an HDInsight Spark cluster using an Azure Resource Manager template. The template can be found in GitHub.

References: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing>

### **QUESTION 5**

**DRAG DROP**

Your company plans to create an event processing engine to handle streaming data from Twitter.

The data engineering team uses Azure Event Hubs to ingest the streaming data.

You need to implement a solution that uses Azure Databricks to receive the streaming data from the Azure Event Hubs.

Which three actions should you recommend be performed 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 and configure a Notebook that consumes the streaming data.

Import data from Blob storage.

Use Environment variables to define the Apache Spark connection.

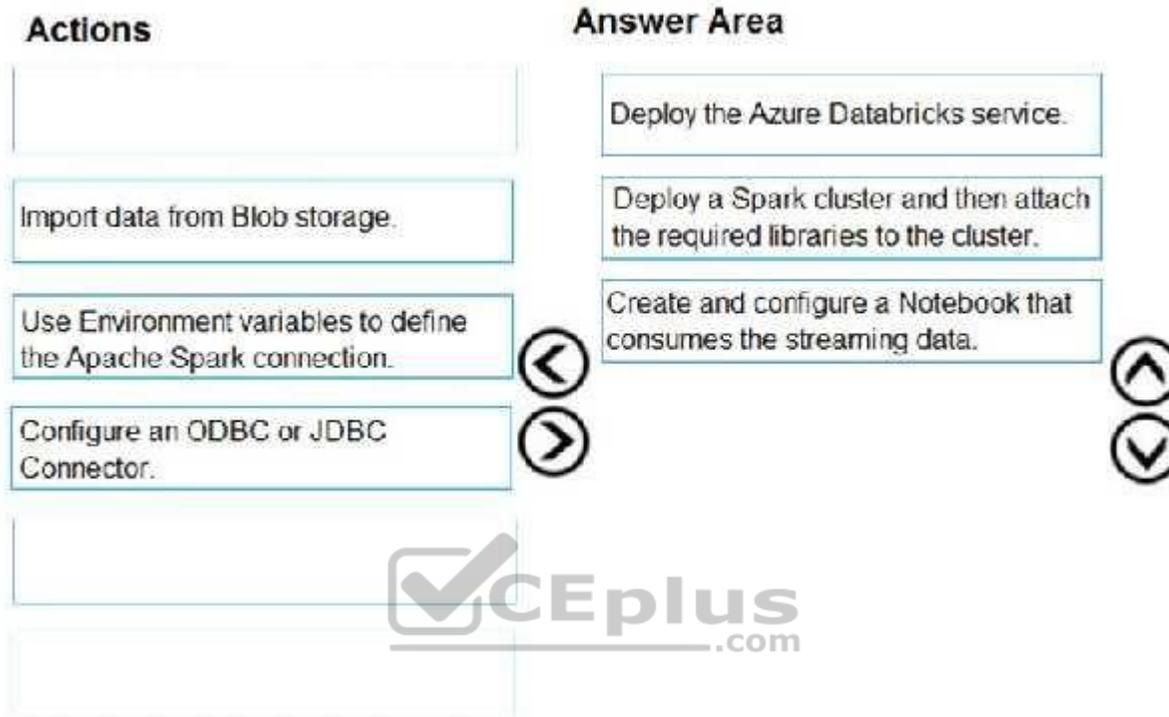
Configure an ODBC or JDBC Connector.

Deploy the Azure Databricks service.

Deploy a Spark cluster and then attach the required libraries to the cluster.



Correct Answer:



**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Step 1: Deploy the Azure Databricks service

Create an Azure Databricks workspace by setting up an Azure Databricks Service.

Step 2: Deploy a Spark cluster and then attach the required libraries to the cluster.

To create a Spark cluster in Databricks, in the Azure portal, go to the Databricks workspace that you created, and then select Launch Workspace.

Attach libraries to Spark cluster: you use the Twitter APIs to send tweets to Event Hubs. You also use the Apache Spark Event Hubs connector to read and write data into Azure Event Hubs. To use these APIs as part of your cluster, add them as libraries to Azure Databricks and associate them with your Spark cluster.

Step 3: Create and configure a Notebook that consumes the streaming data.

<https://vceplus.com/>

You create a notebook named ReadTweetsFromEventhub in Databricks workspace. ReadTweetsFromEventHub is a consumer notebook you use to read the tweets from Event Hubs.

References: <https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-stream-from-eventhubs>

#### QUESTION 6

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:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

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

References: <https://docs.databricks.com/user-guide/notebooks/notebook-use.html#mix-languages>

#### QUESTION 7

You use Azure Stream Analytics to receive Twitter data from Azure Event Hubs and to output the data to an Azure Blob storage account.

You need to output the count of tweets during the last five minutes every five minutes. Each tweet must only be counted once.

Which windowing function should you use?

- A. a five-minute Session window
- B. a five-minute Sliding window
- C. a five-minute Tumbling window
- D. a five-minute Hopping window that has one-minute hop



**Correct Answer:** C

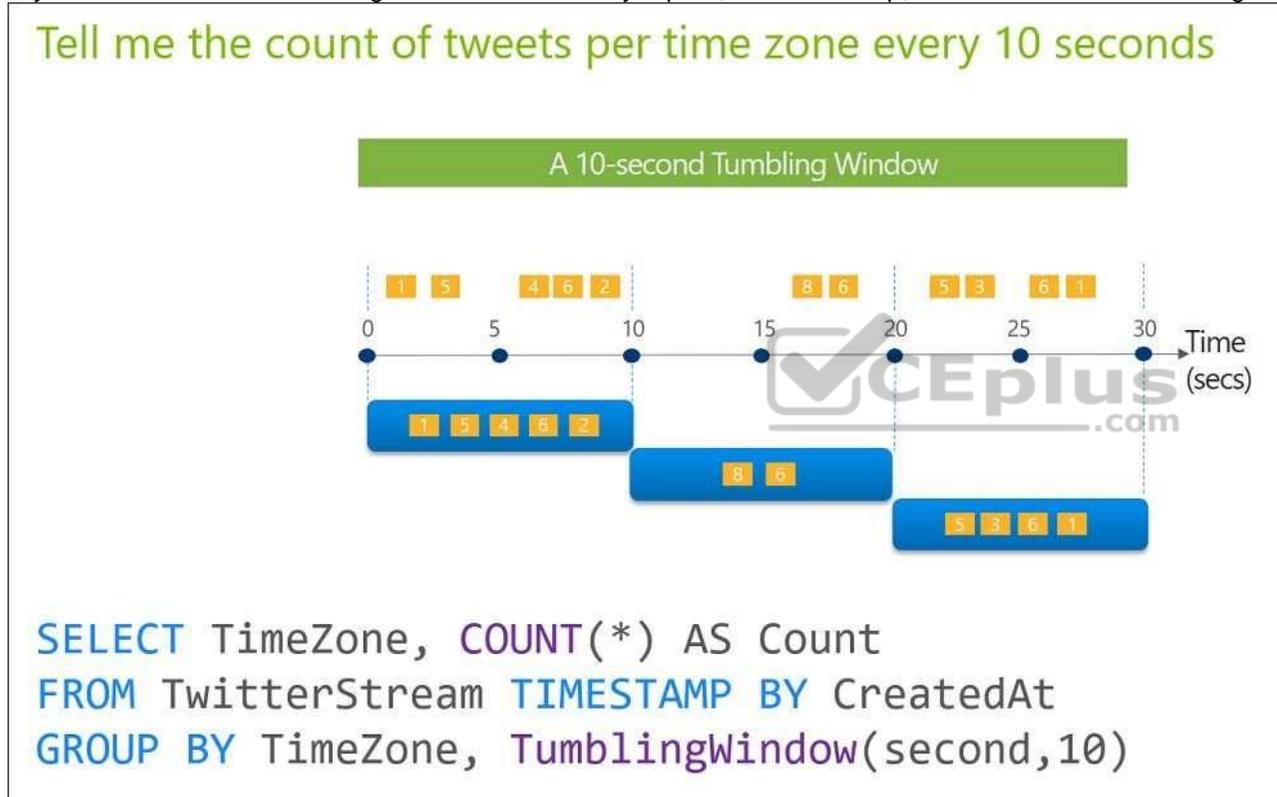
**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

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



References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-window-functions>

**QUESTION 8**

You are developing a solution that will stream to Azure Stream Analytics. The solution will have both streaming data and reference data.

Which input type should you use for the reference data?

<https://vceplus.com/>

- A. Azure Cosmos DB
- B. Azure Event Hubs
- C. Azure Blob storage
- D. Azure IoT Hub

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Stream Analytics supports Azure Blob storage and Azure SQL Database as the storage layer for Reference Data.

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-use-reference-data>

#### QUESTION 9

You have an Azure Storage account and an Azure SQL data warehouse in the UK South region.

You need to copy blob data from the storage account to the data warehouse by using Azure Data Factory.

The solution must meet the following requirements:

- Ensure that the data remains in the UK South region at all times. ▪
- Minimize administrative effort.

Which type of integration runtime should you use?

- A. Azure integration runtime
- B. Self-hosted integration runtime
- C. Azure-SSIS integration runtime

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

IR type	Public network	Private network
Azure	Data Flow Data movement Activity dispatch	
Self-hosted	Data movement Activity dispatch	Data movement Activity dispatch
Azure-SSIS	SSIS package execution	SSIS package execution

Incorrect Answers:

B: Self-hosted integration runtime is to be used On-premises.



References: <https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime>

### QUESTION 10

**Note: This question is part of series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.**

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Use Azure Data Factory to convert the parquet files to CSV files
2. Create an external data source pointing to the Azure storage account
3. Create an external file format and external table using the external data source
4. Load the data using the `INSERT...SELECT` statement

Does the solution meet the goal?

<https://vceplus.com/>

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

There is no need to convert the parquet files to CSV files.

You load the data using the CREATE TABLE AS SELECT statement.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

#### QUESTION 11

**Note: This question is part of series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.**

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Create an external data source pointing to the Azure storage account
2. Create an external file format and external table using the external data source
3. Load the data using the `INSERT...SELECT` statement Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You load the data using the CREATE TABLE AS SELECT statement.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

<https://vceplus.com/>

## QUESTION 12

**Note:** This question is part of series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You develop a data ingestion process that will import data to a Microsoft Azure SQL Data Warehouse. The data to be ingested resides in parquet files stored in an Azure Data Lake Gen 2 storage account.

You need to load the data from the Azure Data Lake Gen 2 storage account into the Azure SQL Data Warehouse.

Solution:

1. Create an external data source pointing to the Azure storage account
2. Create a workload group using the Azure storage account name as the pool name
3. Load the data using the `INSERT...SELECT` statement Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You need to create an external file format and external table using the external data source.

You then load the data using the `CREATE TABLE AS SELECT` statement.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-load-from-azure-data-lake-store>

## QUESTION 13

You develop data engineering solutions for a company.

You must integrate the company's on-premises Microsoft SQL Server data with Microsoft Azure SQL Database. Data must be transformed incrementally.

You need to implement the data integration solution.

Which tool should you use to configure a pipeline to copy data?

- A. Use the Copy Data tool with Blob storage linked service as the source
- B. Use Azure PowerShell with SQL Server linked service as a source
- C. Use Azure Data Factory UI with Blob storage linked service as a source

<https://vceplus.com/>

D. Use the .NET Data Factory API with Blob storage linked service as the source

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

The Integration Runtime is a customer managed data integration infrastructure used by Azure Data Factory to provide data integration capabilities across different network environments.

A linked service defines the information needed for Azure Data Factory to connect to a data resource. We have three resources in this scenario for which linked services are needed:

- On-premises SQL Server
- Azure Blob Storage
- Azure SQL database

Note: Azure Data Factory is a fully managed cloud-based data integration service that orchestrates and automates the movement and transformation of data. The key concept in the ADF model is pipeline. A pipeline is a logical grouping of Activities, each of which defines the actions to perform on the data contained in Datasets. Linked services are used to define the information needed for Data Factory to connect to the data resources.

References: <https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/move-sql-azure-ADF>

## QUESTION 14

### HOTSPOT

A company runs Microsoft Dynamics CRM with Microsoft SQL Server on-premises. SQL Server Integration Services (SSIS) packages extract data from Dynamics CRM APIs, and load the data into a SQL Server data warehouse.

The datacenter is running out of capacity. Because of the network configuration, you must extract on premises data to the cloud over https. You cannot open any additional ports. The solution must implement the least amount of effort.

You need to create the pipeline system.

Which component should you use? To answer, select the appropriate technology in the dialog box in the answer area.

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

**Hot Area:**

Action	Technology
Extract SQL data on-premises	Self-hosted integration runtime
	Azure-SSIS integration runtime
	Azure integration runtime
	Source
Load SQL data warehouse	Self-hosted integration runtime
	Azure-SSIS integration runtime
	Azure integration runtime
	Sink

Correct Answer:



Action	Technology
Extract SQL data on-premises	Self-hosted integration runtime
	Azure-SSIS integration runtime
	Azure integration runtime
	Source
Load SQL data warehouse	Self-hosted integration runtime
	Azure-SSIS integration runtime
	Azure integration runtime
	Sink

Section: [none]  
Explanation

**Explanation/Reference:**

Explanation:

Box 1: Source

For Copy activity, it requires source and sink linked services to define the direction of data flow.

Copying between a cloud data source and a data source in private network: if either source or sink linked service points to a self-hosted IR, the copy activity is executed on that self-hosted Integration Runtime.

Box 2: Self-hosted integration runtime

A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network, and it can dispatch transform activities against compute resources in an on-premises network or an Azure virtual network. The installation of a self-hosted integration runtime needs on an on-premises machine or a virtual machine (VM) inside a private network.

References: <https://docs.microsoft.com/en-us/azure/data-factory/create-self-hosted-integration-runtime>

**QUESTION 15**

You develop data engineering solutions for a company.

You need to ingest and visualize real-time Twitter data by using Microsoft Azure.

Which three technologies should you use? Each correct answer presents part of the solution.

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

- A. Event Grid topic
- B. Azure Stream Analytics Job that queries Twitter data from an Event Hub
- C. Azure Stream Analytics Job that queries Twitter data from an Event Grid
- D. Logic App that sends Twitter posts which have target keywords to Azure
- E. Event Grid subscription
- F. Event Hub instance

**Correct Answer:** BDF

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You can use Azure Logic apps to send tweets to an event hub and then use a Stream Analytics job to read from event hub and send them to PowerBI.

References: <https://community.powerbi.com/t5/Integrations-with-Files-and/Twitter-streaming-analytics-step-by-step/td-p/9594>

<https://vceplus.com/>

## QUESTION 16

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You plan to create an Azure Databricks workspace that has a tiered structure. The workspace will contain the following three workloads:

- A workload for data engineers who will use Python and SQL
- A workload for jobs that will run notebooks that use Python, Spark, Scala, and SQL
- A workload that data scientists will use to perform ad hoc analysis in Scala and R

The enterprise architecture team at your company identifies the following standards for Databricks environments:

- The data engineers must share a cluster.
- The job cluster will be managed by using a request process whereby data scientists and data engineers provide packaged notebooks for deployment to the cluster.
- All the data scientists must be assigned their own cluster that terminates automatically after 120 minutes of inactivity. Currently, there are three data scientists.

You need to create the Databrick clusters for the workloads.

Solution: You create a Standard cluster for each data scientist, a High Concurrency cluster for the data engineers, and a Standard cluster for the jobs.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

We would need a High Concurrency cluster for the jobs.

Note:

Standard clusters are recommended for a single user. Standard can run workloads developed in any language: Python, R, Scala, and SQL.

A high concurrency cluster is a managed cloud resource. The key benefits of high concurrency clusters are that they provide Apache Spark-native fine-grained sharing for maximum resource utilization and minimum query latencies.

<https://vceplus.com/>

References:

<https://docs.azuredatabricks.net/clusters/configure.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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You plan to create an Azure Databricks workspace that has a tiered structure. The workspace will contain the following three workloads:

- A workload for data engineers who will use Python and SQL
- A workload for jobs that will run notebooks that use Python, Spark, Scala, and SQL
- A workload that data scientists will use to perform ad hoc analysis in Scala and R

The enterprise architecture team at your company identifies the following standards for Databricks environments:

- The data engineers must share a cluster.
- The job cluster will be managed by using a request process whereby data scientists and data engineers provide packaged notebooks for deployment to the cluster.
- All the data scientists must be assigned their own cluster that terminates automatically after 120 minutes of inactivity. Currently, there are three data scientists.

You need to create the Databrick clusters for the workloads.

Solution: You create a Standard cluster for each data scientist, a High Concurrency cluster for the data engineers, and a High Concurrency cluster for the jobs.

Does this meet the goal?



<https://vceplus.com/>

- A. Yes
- B. No

<https://vceplus.com/>

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

We need a High Concurrency cluster for the data engineers and the jobs.

Note:

Standard clusters are recommended for a single user. Standard can run workloads developed in any language: Python, R, Scala, and SQL.

A high concurrency cluster is a managed cloud resource. The key benefits of high concurrency clusters are that they provide Apache Spark-native fine-grained sharing for maximum resource utilization and minimum query latencies.

References:

<https://docs.azuredatabricks.net/clusters/configure.html>

#### **QUESTION 18**

You have an Azure Stream Analytics query. The query returns a result set that contains 10,000 distinct values for a column named clusterID.

You monitor the Stream Analytics job and discover high latency.

You need to reduce the latency.

Which two actions should you perform? Each correct answer presents a complete solution.

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

- A. Add a pass-through query.
- B. Add a temporal analytic function.
- C. Scale out the query by using `PARTITION BY`.
- D. Convert the query to a reference query.
- E. Increase the number of streaming units.

**Correct Answer:** CE

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

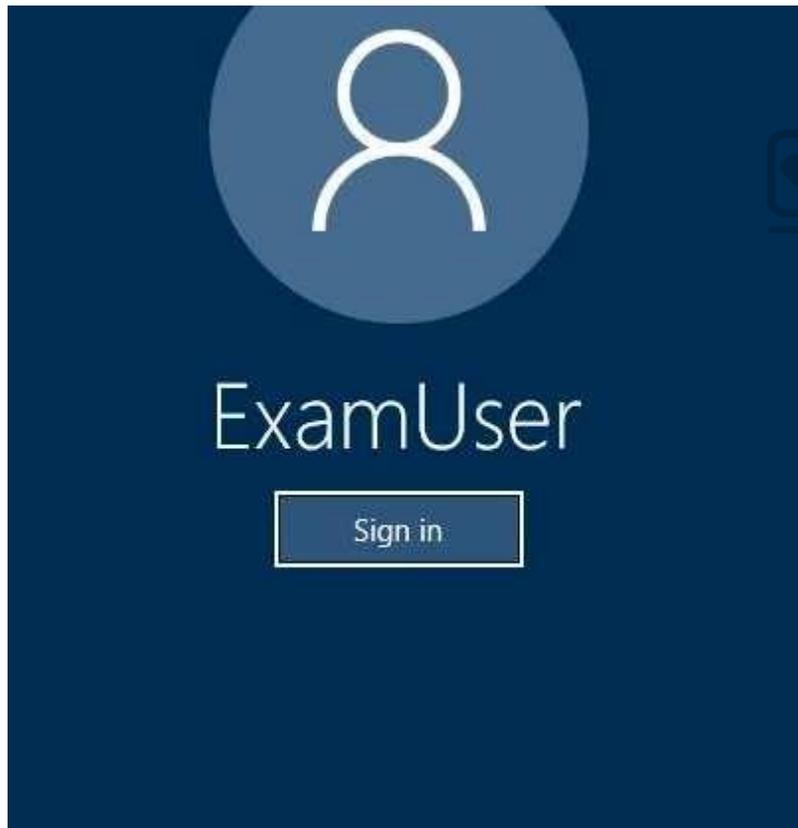
<https://vceplus.com/>

C: Scaling a Stream Analytics job takes advantage of partitions in the input or output. Partitioning lets you divide data into subsets based on a partition key. A process that consumes the data (such as a Streaming Analytics job) can consume and write different partitions in parallel, which increases throughput.

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

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

**QUESTION 19**  
SIMULATION



<https://vceplus.com/>

Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10277521

You plan to create multiple pipelines in a new Azure Data Factory V2.

You need to create the data factory, and then create a scheduled trigger for the planned pipelines. The trigger must execute every two hours starting at 24:00:00.

To complete this task, sign in to the Azure portal.

**Correct Answer:** See the explanation below.

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Step 1: Create a new Azure Data Factory V2

1. Go to the Azure portal.

2. Select Create a resource on the left menu, select Analytics, and then select Data Factory.



Home > New

## New

Search the Marketplace

Azure Marketplace [See all](#)      Featured [See all](#)

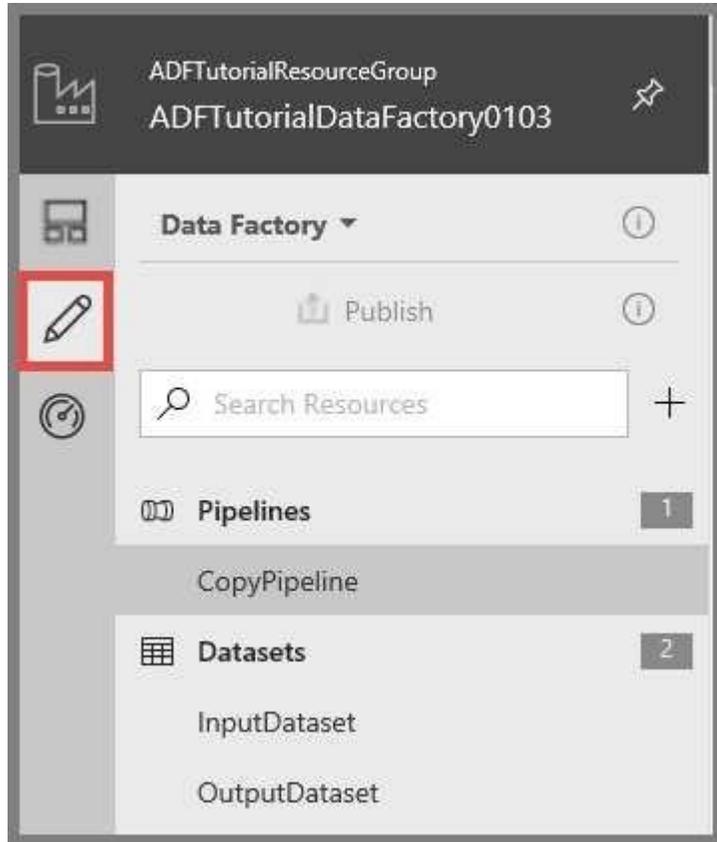
Get started		Azure Data Explorer <a href="#">Learn more</a>
Recently created		
AI + Machine Learning		HDInsight <a href="#">Quickstart tutorial</a>
<b>Analytics</b>		
Blockchain		Data Lake Analytics <a href="#">Quickstart tutorial</a>
Compute		
Containers		Stream Analytics job <a href="#">Quickstart tutorial</a>
Databases		
Developer Tools		Analysis Services <a href="#">Quickstart tutorial</a>
DevOps		
Identity		Azure Databricks <a href="#">Quickstart tutorial</a>
Integration		
Internet of Things		Power BI Embedded <a href="#">Quickstart tutorial</a>
Media		
Mixed Reality		SQL Data Warehouse <a href="#">Quickstart tutorial</a>
IT & Management Tools		
Networking		Data Lake Storage Gen1 <a href="#">Quickstart tutorial</a>
Software as a Service (SaaS)		

4. On the New data factory page, enter a name.
5. For Subscription, select your Azure subscription in which you want to create the data factory.
6. For Resource Group, use one of the following steps:
  - Select Use existing, and select an existing resource group from the list. ▪Select Create new, and enter the name of a resource group.
7. For Version, select V2.
8. For Location, select the location for the data factory.
9. Select Create.
10. After the creation is complete, you see the Data Factory page.

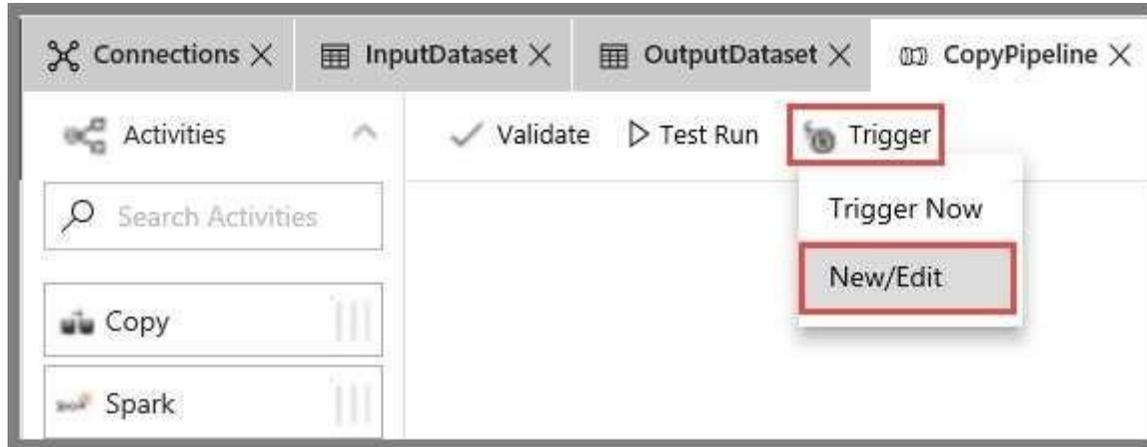
#### Step 2: Create a schedule trigger for the Data Factory

1. Select the Data Factory you created, and switch to the Edit tab.

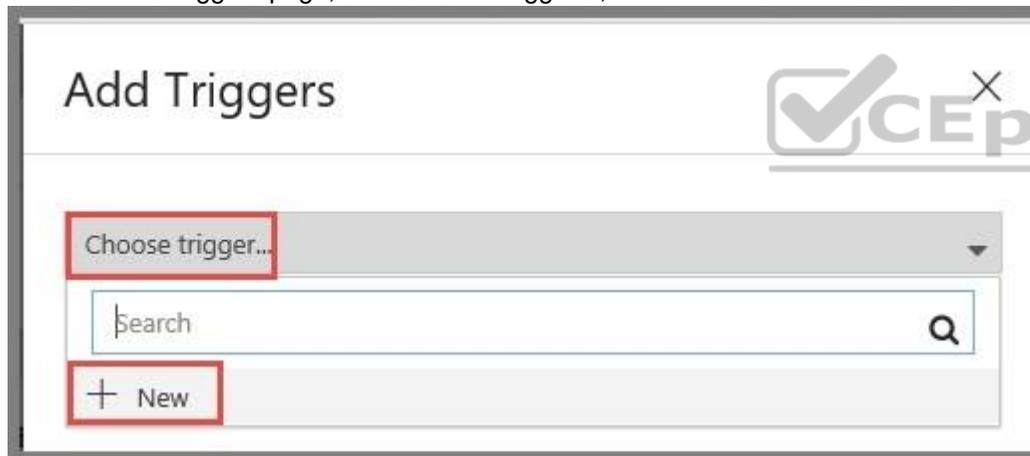




2. Click Trigger on the menu, and click New/Edit.



3. In the Add Triggers page, click Choose trigger..., and click New.



4. In the New Trigger page, do the following steps:

- Confirm that Schedule is selected for Type.
- Specify the start datetime of the trigger for Start Date (UTC) to: 24:00:00
- Specify Recurrence for the trigger. Select Every Hour, and enter 2 in the text box.

Type \*

Schedule  Tumbling Window

Start Date (UTC) \*

01/03/2018, 10:54 PM

Recurrence \*

Every Minute

Every

1

Minute(s)

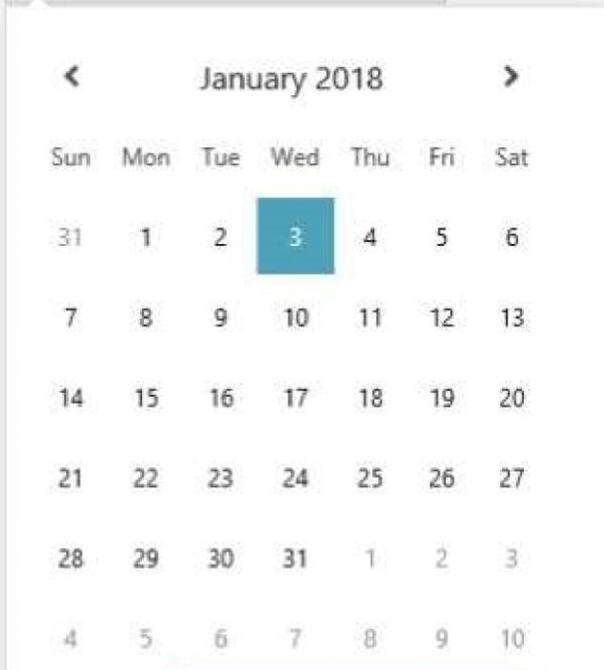
End \*

1

On Date

End On (UTC) \*

01/03/2018, 11:10 PM



January 2018

Sun	Mon	Tue	Wed	Thu	Fri	Sat
31	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31	1	2	3
4	5	6	7	8	9	10



Next

En

2

11

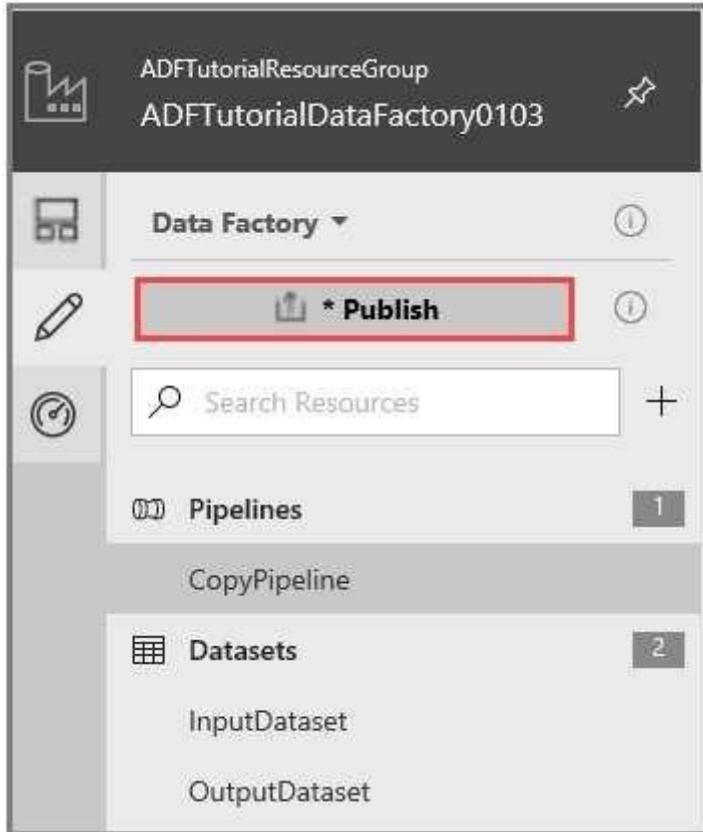
:

10

PM

3

5. In the New Trigger window, check the Activated option, and click Next.
6. In the New Trigger page, review the warning message, and click Finish.
7. Click Publish to publish changes to Data Factory. Until you publish changes to Data Factory, the trigger does not start triggering the pipeline runs.



References:

<https://docs.microsoft.com/en-us/azure/data-factory/quickstart-create-data-factory-portal>

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-schedule-trigger> **Testlet 2**

## Overview

## Current environment

<https://vceplus.com/>

Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.

The majority of the company's data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

Applications	Tier	Replication	Notes
Internal Contoso	1	Yes	
Internal Contoso	2	SQL Data Sync	Data Sync between databases
Internal Partner	3	Yes	Replicate to Partner
External Contoso	4,5,6	Yes	
External Partner	7,8	No	Partner managed
Internal Distribution and Sales	9	Yes, once ingested at branches	Data ingested from Contoso branches
External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources

The company has a reporting infrastructure that ingests data from local databases and partner services. Partners services consists of distributors, wholesales, and retailers across the world. The company performs daily, weekly, and monthly reporting.

### Requirements

Tier 3 and Tier 6 through Tier 8 application must use database density on the same server and Elastic pools in a cost-effective manner.

Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.

A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of server going offline.

Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.

- Tier 1 internal applications on the premium P2 tier ▪
  - Tier 2 internal applications on the standard S4 tier
- The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.

Tier 7 and Tier 8 partner access must be restricted to the database only.

In addition to default Azure backup behavior, Tier 4 and 5 databases must be on a backup strategy that performs a transaction log backup every hour, a differential backup of databases every day and a full back up every week.

Back up strategies must be put in place for all other standalone Azure SQL Databases using Azure SQL-provided backup storage and capabilities.

### Databases

Contoso requires their data estate to be designed and implemented in the Azure Cloud. Moving to the cloud must not inhibit access to or availability of data.

Databases:

Tier 1 Database must implement data masking using the following masking logic:

Data type	Masking requirement
A	Mask 4 or less string data type characters
B	Mask first letter and domain
C	Mask everything except characters at the beginning and end

Tier 2 databases must sync between branches and cloud databases and in the event of conflicts must be set up for conflicts to be won by on-premises databases.

Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.

Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.

A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of a server going offline.

Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.

- Tier 1 internal applications on the premium P2 tier
- Tier 2 internal applications on the standard S4 tier

### Reporting

### Security and monitoring

### Security

A method of managing multiple databases in the cloud at the same time is must be implemented to streamlining data management and limiting management access to only those requiring access.

### Monitoring

Monitoring must be set up on every database. Contoso and partners must receive performance reports as part of contractual agreements.

Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.

The Azure SQL Data Warehouse cache must be monitored when the database is being used. A dashboard monitoring key performance indicators (KPIs) indicated by traffic lights must be created and displayed based on the following metrics:

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

Existing Data Protection and Security compliances require that all certificates and keys are internally managed in an on-premises storage.

You identify the following reporting requirements:

- Azure Data Warehouse must be used to gather and query data from multiple internal and external databases
- Azure Data Warehouse must be optimized to use data from a cache
- Reporting data aggregated for external partners must be stored in Azure Storage and be made available during regular business hours in the connecting regions
- Reporting strategies must be improved to real time or near real time reporting cadence to improve competitiveness and the general supply chain
- Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company's main office
- Tier 10 reporting data must be stored in Azure Blobs

### Issues

Team members identify the following issues:

- Both internal and external client application run complex joins, equality searches and group-by clauses. Because some systems are managed externally, the queries will not be changed or optimized by Contoso
- External partner organization data formats, types and schemas are controlled by the partner companies
- Internal and external database development staff resources are primarily SQL developers familiar with the Transact-SQL language.
- Size and amount of data has led to applications and reporting solutions not performing are required speeds
- Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
- The company maintains several legacy client applications. Data for these applications remains isolated form other applications. This has led to hundreds of databases being provisioned on a per application basis

<https://vceplus.com/>

### QUESTION 1

You need to process and query ingested Tier 9 data.

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

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

- A. Azure Notification Hub
- B. Transact-SQL statements
- C. Azure Cache for Redis
- D. Apache Kafka statements
- E. Azure Event Grid
- F. Azure Stream Analytics

**Correct Answer:** EF

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Event Hubs provides a Kafka endpoint that can be used by your existing Kafka based applications as an alternative to running your own Kafka cluster.

You can stream data into Kafka-enabled Event Hubs and process it with Azure Stream Analytics, in the following steps:

- Create a Kafka enabled Event Hubs namespace.
- Create a Kafka client that sends messages to the event hub.
- Create a Stream Analytics job that copies data from the event hub into an Azure blob storage.

Scenario:

Internal Distribution and Sales	9	Yes, once ingested at branches	Data ingested from Contoso branches
---------------------------------	---	--------------------------------	-------------------------------------

Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company's main office

References:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-kafka-stream-analytics>

### QUESTION 2

You need to set up Azure Data Factory pipelines to meet data movement requirements.

Which integration runtime should you use?

- A. self-hosted integration runtime
- B. Azure-SSIS Integration Runtime
- C. .NET Common Language Runtime (CLR)
- D. Azure integration runtime

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

The following table describes the capabilities and network support for each of the integration runtime types:

IR type	Public network	Private network
Azure	Data movement Activity dispatch	
Self-hosted	Data movement Activity dispatch	Data movement Activity dispatch
Azure-SSIS	SSIS package execution	SSIS package execution

Scenario: The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.

References:

<https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime>

## Question Set 1

### QUESTION 1

You plan to use Microsoft Azure SQL Database instances with strict user access control. A user object must:

- Move with the database if it is run elsewhere
- Be able to create additional users

You need to create the user object with correct permissions.

Which two Transact-SQL commands should you run? Each correct answer presents part of the solution.

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

- A. ALTER LOGIN Mary WITH PASSWORD = 'strong\_password';
- B. CREATE LOGIN Mary WITH PASSWORD = 'strong\_password';
- C. ALTER ROLE db\_owner ADD MEMBER Mary;
- D. CREATE USER Mary WITH PASSWORD = 'strong\_password';
- E. GRANT ALTER ANY USER TO Mary;

**Correct Answer:** CD

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

C: ALTER ROLE adds or removes members to or from a database role, or changes the name of a user-defined database role.

Members of the db\_owner fixed database role can perform all configuration and maintenance activities on the database, and can also drop the database in SQL Server.

D: CREATE USER adds a user to the current database.

Note: Logins are created at the server level, while users are created at the database level. In other words, a login allows you to connect to the SQL Server service (also called an instance), and permissions inside the database are granted to the database users, not the logins. The logins will be assigned to server roles (for example, serveradmin) and the database users will be assigned to roles within that database (eg. db\_datareader, db\_backupoperator).

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

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

**QUESTION 2 HOTSPOT**

<https://vceplus.com/>

Your company uses Azure SQL Database and Azure Blob storage.

All data at rest must be encrypted by using the company's own key. The solution must minimize administrative effort and the impact to applications which use the database.

You need to configure security.

What should you implement? To answer, select the appropriate option in the answer area.

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

**Hot Area:**



## Answer Area

Service	Encryption at rest
Azure SQL Database	<ul style="list-style-type: none"><li>always encrypted</li><li>cell-level encryption</li><li>row-level security</li><li>transparent data encryption</li></ul>
Azure Storage	<ul style="list-style-type: none"><li>Azure disk encryption</li><li>secure transport layer security (TLS)</li><li>storage account keys</li><li>default storage service encryption</li></ul>

Correct Answer:

## Answer Area

Service	Encryption at rest
Azure SQL Database	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <ul style="list-style-type: none"> <li>always encrypted</li> <li>cell-level encryption</li> <li>row-level security</li> <li style="background-color: #d9ead3;">transparent data encryption</li> </ul> </div>
Azure Storage	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <ul style="list-style-type: none"> <li>Azure disk encryption</li> <li>secure transport layer security (TLS)</li> <li style="background-color: #d9ead3;">storage account keys</li> <li>default storage service encryption</li> </ul> </div>

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Box 1: transparent data encryption

TDE with customer-managed keys in Azure Key Vault allows to encrypt the Database Encryption Key (DEK) with a customer-managed asymmetric key called TDE Protector. This is also generally referred to as Bring Your Own Key (BYOK) support for Transparent Data Encryption.

Note: Transparent data encryption encrypts the storage of an entire database by using a symmetric key called the database encryption key. This database encryption key is protected by the transparent data encryption protector.

<https://vceplus.com/>

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Data Warehouse 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.

Box 2: Storage account keys

You can rely on Microsoft-managed keys for the encryption of your storage account, or you can manage encryption with your own keys, together with Azure Key Vault.

References: <https://docs.microsoft.com/en-us/azure/sql-database/transparent-data-encryption-azure-sql> <https://docs.microsoft.com/en-us/azure/storage/common/storage-service-encryption>

### QUESTION 3

You develop data engineering solutions for a company.

A project requires the deployment of data to Azure Data Lake Storage.

You need to implement role-based access control (RBAC) so that project members can manage the Azure Data Lake Storage resources.

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

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

- A. Assign Azure AD security groups to Azure Data Lake Storage.
- B. Configure end-user authentication for the Azure Data Lake Storage account.
- C. Configure service-to-service authentication for the Azure Data Lake Storage account.
- D. Create security groups in Azure Active Directory (Azure AD) and add project members.
- E. Configure access control lists (ACL) for the Azure Data Lake Storage account.

**Correct Answer:** ADE

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

AD: Create security groups in Azure Active Directory. Assign users or security groups to Data Lake Storage Gen1 accounts.

E: Assign users or security groups as ACLs to the Data Lake Storage Gen1 file system

References: <https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-secure-data>

**QUESTION 4**  
SIMULATION



Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10277521

<https://vceplus.com/>

You need to ensure that an email notification is sent to admin@contoso.com if a suspicious login to an Azure SQL database named db2 is detected.

To complete this task, sign in to the Azure portal.

**Correct Answer:** See the explanation below.

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Set up Advanced Threat Protection in the Azure portal.

1. From the Azure portal navigate to the configuration page of the Azure SQL Database db2, which you want to protect. In the security settings, select Advanced Data Security.
2. On the Advanced Data Security configuration page:

Enable Advanced Data Security on the server.

In Advanced Threat Protection Settings, in the Send alerts to text box, enter admin@contoso.com to receive security alerts upon detection of anomalous database activities.



### vanazuressqlserver - Advanced Data Security

SQL server

Search (Ctrl+/)

Save Discard Feedback

#### ADVANCED DATA SECURITY

ON  OFF

Advanced Data Security costs 15 USD/server/month. It includes Data Discovery & Classification, Vulnerability Assessment and Advanced Threat Protection. We invite you to a trial period for the first 30 days, without charge.

#### VULNERABILITY ASSESSMENT SETTINGS

Subscription >

SQL DB Content >

Storage account >

Periodic recurring scans  ON  OFF

Send scan reports to

Also send email notification to admins and subscription owners

#### ADVANCED THREAT PROTECTION SETTINGS

Send alerts to

Email addresses ✓

Also send email notification to admins and subscription owners

Advanced Threat Protection types >

All >

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Settings

Quick start

Failover groups

Manage Backups

Active Directory admin

SQL databases

SQL elastic pools

Deleted databases

Import/Export history

DTU quota

Properties

Locks

Export template

Security

Advanced Data Security

References:

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



## Testlet 2

### Overview

#### Current environment

Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.

The majority of the company's data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

Applications	Tier	Replication	Notes
Internal Contoso	1	Yes	
Internal Contoso	2	SQL Data Sync	Data Sync between databases
Internal Partner	3	Yes	Replicate to Partner
External Contoso	4,5,6	Yes	
External Partner	7,8	No	Partner managed
Internal Distribution and Sales	9	Yes, once ingested at branches	Data ingested from Contoso branches
External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources

The company has a reporting infrastructure that ingests data from local databases and partner services. Partners services consists of distributors, wholesales, and retailers across the world. The company performs daily, weekly, and monthly reporting.

#### Requirements

Tier 3 and Tier 6 through Tier 8 application must use database density on the same server and Elastic pools in a cost-effective manner.

Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.

A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of server going offline.

Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.

- Tier 1 internal applications on the premium P2 tier
  - Tier 2 internal applications on the standard S4 tier
- The solution must support migrating databases that support external and internal application to Azure SQL

<https://vceplus.com/>

Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.

Tier 7 and Tier 8 partner access must be restricted to the database only.

In addition to default Azure backup behavior, Tier 4 and 5 databases must be on a backup strategy that performs a transaction log backup every hour, a differential backup of databases every day and a full back up every week.

Back up strategies must be put in place for all other standalone Azure SQL Databases using Azure SQL-provided backup storage and capabilities.

### Databases

Contoso requires their data estate to be designed and implemented in the Azure Cloud. Moving to the cloud must not inhibit access to or availability of data.

Databases:

Tier 1 Database must implement data masking using the following masking logic:

Data type	Masking requirement
A	Mask 4 or less string data type characters
B	Mask first letter and domain
C	Mask everything except characters at the beginning and end

Tier 2 databases must sync between branches and cloud databases and in the event of conflicts must be set up for conflicts to be won by on-premises databases.

Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.

Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.

A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of a server going offline.

Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.

- Tier 1 internal applications on the premium P2 tier
- Tier 2 internal applications on the standard S4 tier

<https://vceplus.com/>

## Reporting

### Security and monitoring

#### Security

A method of managing multiple databases in the cloud at the same time must be implemented to streamline data management and limit management access to only those requiring access.

#### Monitoring

Monitoring must be set up on every database. Contoso and partners must receive performance reports as part of contractual agreements.

Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.

The Azure SQL Data Warehouse cache must be monitored when the database is being used. A dashboard monitoring key performance indicators (KPIs) indicated by traffic lights must be created and displayed based on the following metrics:

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

Existing Data Protection and Security compliances require that all certificates and keys are internally managed in an on-premises storage.

You identify the following reporting requirements:

- Azure Data Warehouse must be used to gather and query data from multiple internal and external databases
- Azure Data Warehouse must be optimized to use data from a cache
- Reporting data aggregated for external partners must be stored in Azure Storage and be made available during regular business hours in the connecting regions
- Reporting strategies must be improved to real time or near real time reporting cadence to improve competitiveness and the general supply chain
- Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company's main office
- Tier 10 reporting data must be stored in Azure Blobs

#### Issues

Team members identify the following issues:

- Both internal and external client application run complex joins, equality searches and group-by clauses. Because some systems are managed externally, the queries will not be changed or optimized by Contoso

- External partner organization data formats, types and schemas are controlled by the partner companies
- Internal and external database development staff resources are primarily SQL developers familiar with the Transact-SQL language.
- Size and amount of data has led to applications and reporting solutions not performing at required speeds
- Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
- The company maintains several legacy client applications. Data for these applications remains isolated from other applications. This has led to hundreds of databases being provisioned on a per application basis

### QUESTION 1

**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 questions 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 need to configure data encryption for external applications.

Solution:

1. Access the Always Encrypted Wizard in SQL Server Management Studio
2. Select the column to be encrypted
3. Set the encryption type to Randomized
4. Configure the master key to use the Windows Certificate Store
5. Validate configuration results and deploy the solution Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

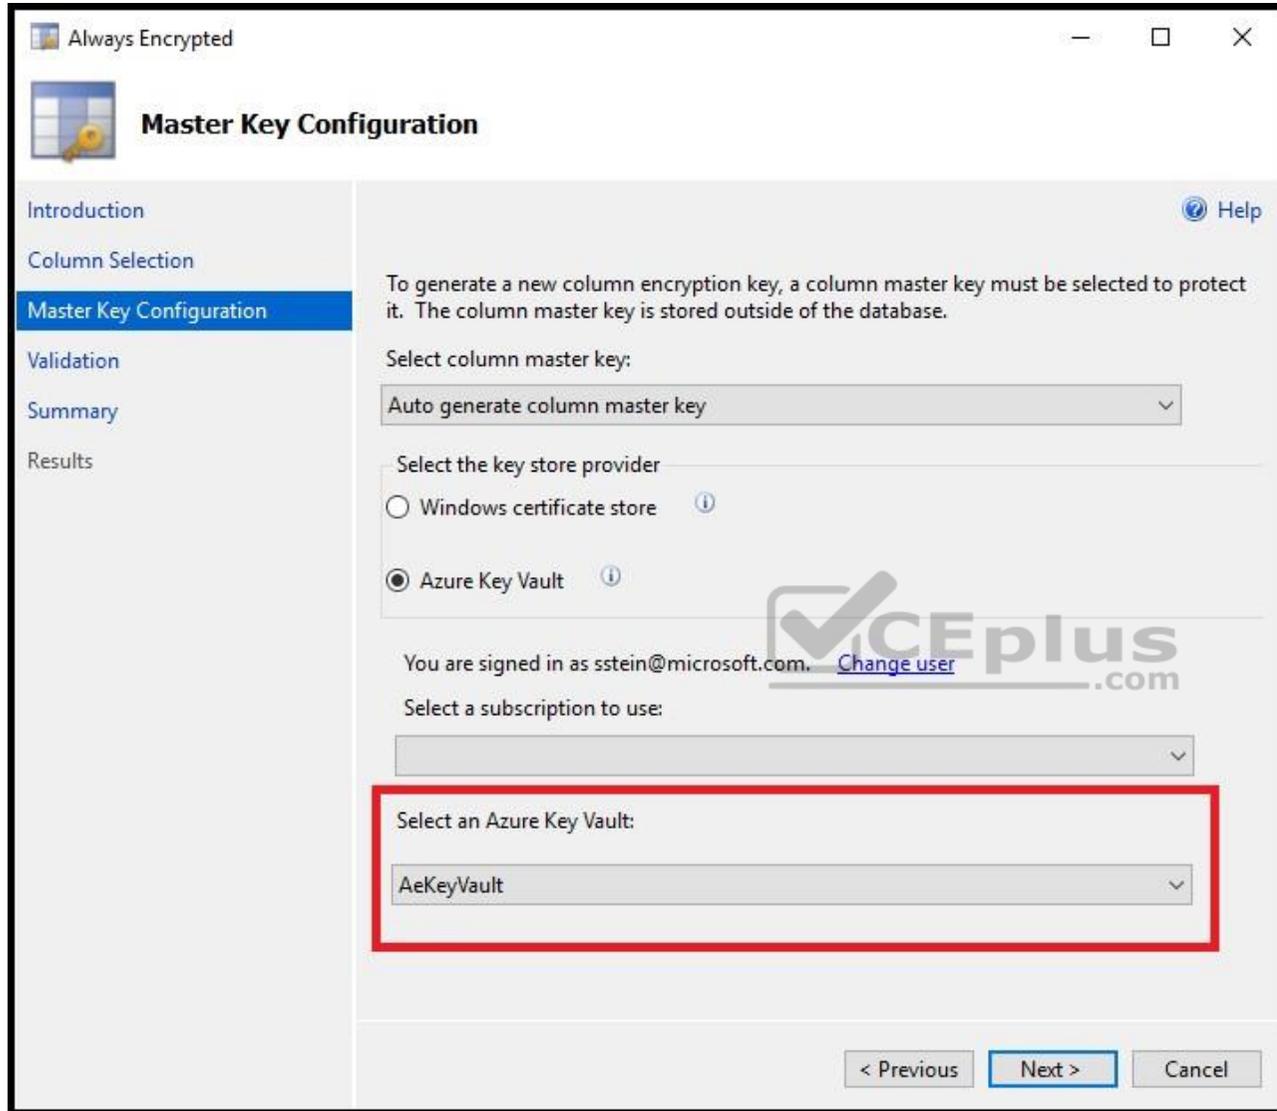
**Explanation**

**Explanation/Reference:**

Explanation:

Use the Azure Key Vault, not the Windows Certificate Store, to store the master key.

Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored. Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).



References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault>

## QUESTION 2

<https://vceplus.com/>

**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 questions 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 need to configure data encryption for external applications.

Solution:

1. Access the Always Encrypted Wizard in SQL Server Management Studio
2. Select the column to be encrypted
3. Set the encryption type to Deterministic
4. Configure the master key to use the Windows Certificate Store
5. Validate configuration results and deploy the solution Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

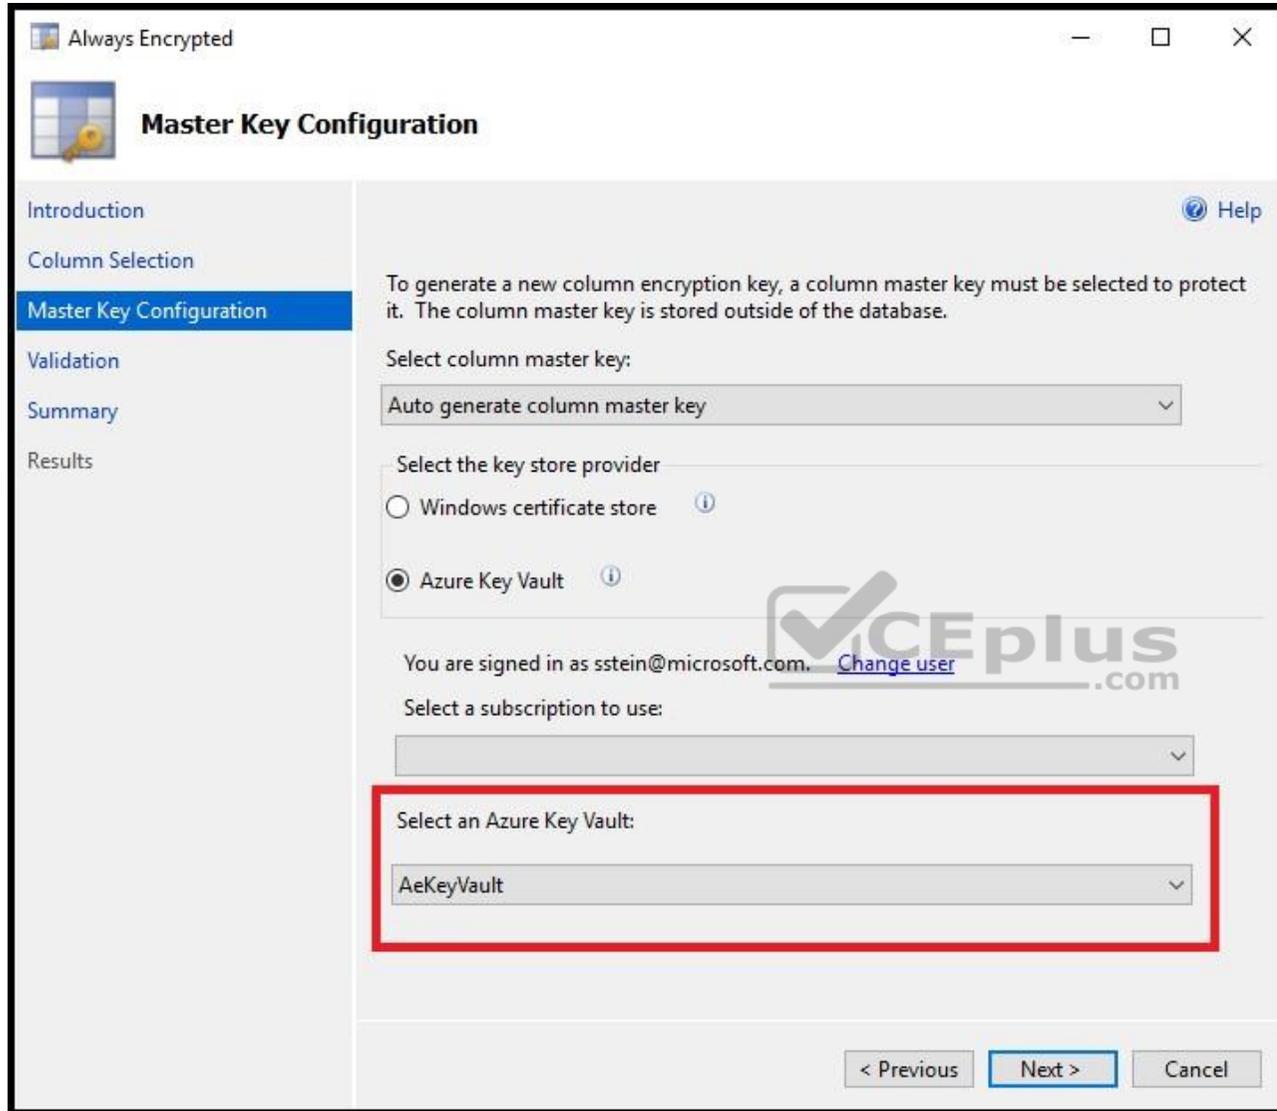


**Explanation/Reference:**

Explanation:

Use the Azure Key Vault, not the Windows Certificate Store, to store the master key.

Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored. Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).



References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault>

### QUESTION 3

<https://vceplus.com/>

**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 questions 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 need to configure data encryption for external applications.

Solution:

1. Access the Always Encrypted Wizard in SQL Server Management Studio
  2. Select the column to be encrypted
  3. Set the encryption type to Deterministic
  4. Configure the master key to use the Azure Key Vault 5. Validate configuration results and deploy the solution Does the solution meet the goal?
- A. Yes  
B. No

**Correct Answer: A**

**Section: [none]**

**Explanation**

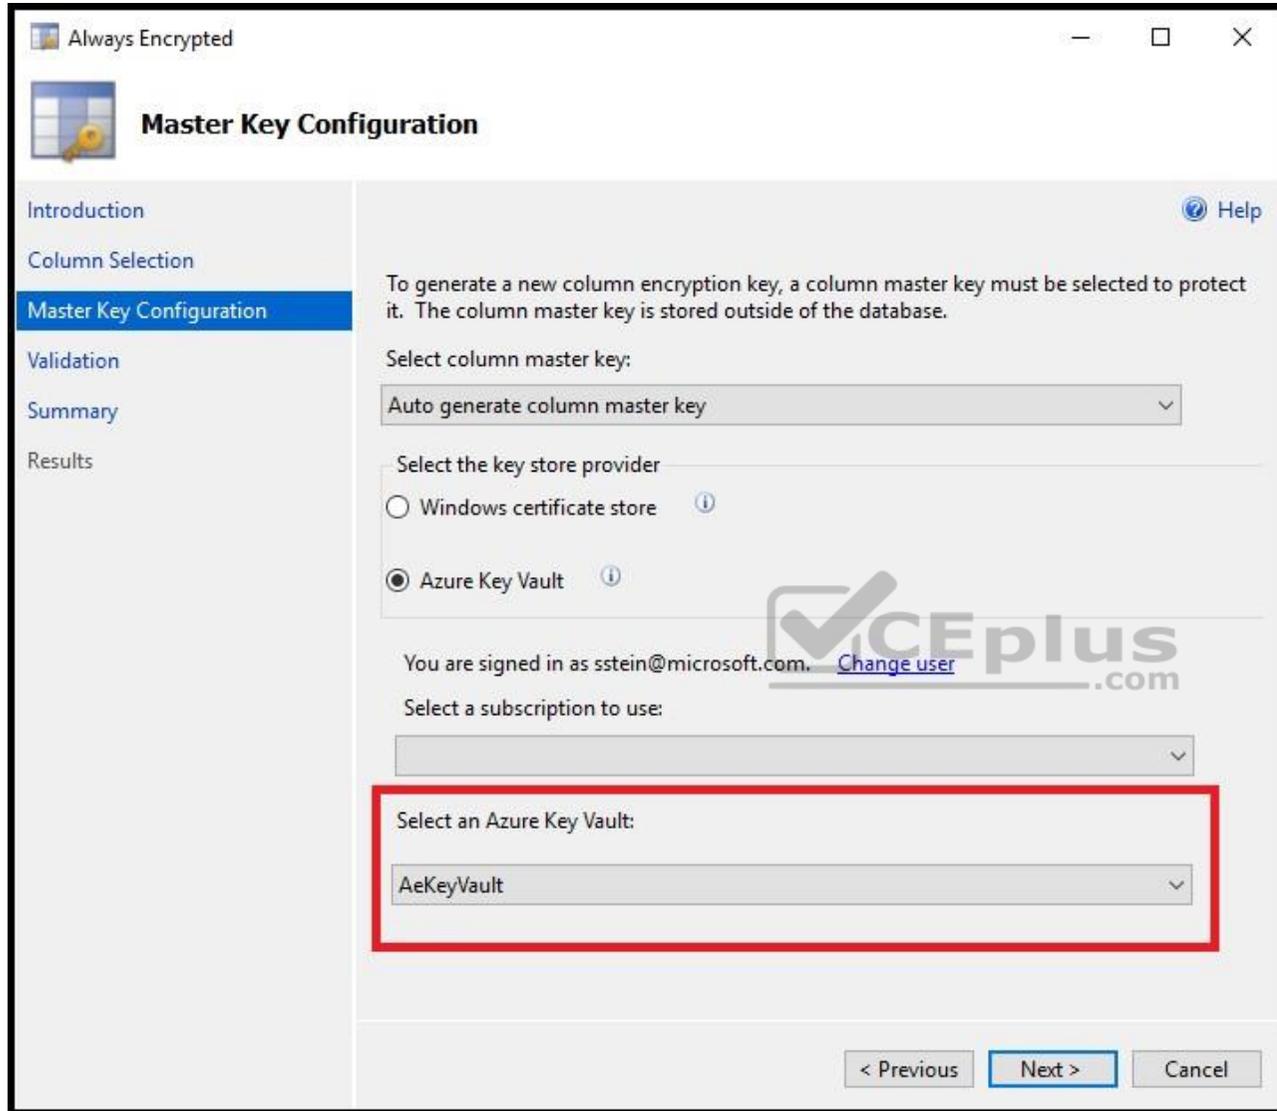
**Explanation/Reference:**

Explanation:

We use the Azure Key Vault, not the Windows Certificate Store, to store the master key.

Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored. Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).





References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault> **Question Set 1**

### QUESTION 1

A company has a Microsoft Azure HDInsight solution that uses different cluster types to process and analyze data. Operations are continuous.

<https://vceplus.com/>

Reports indicate slowdowns during a specific time window.

You need to determine a monitoring solution to track down the issue in the least amount of time.

What should you use?

- A. Azure Log Analytics log search query
- B. Ambari REST API
- C. Azure Monitor Metrics
- D. HDInsight .NET SDK
- E. Azure Log Analytics alert rule query

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Ambari is the recommended tool for monitoring the health for any given HDInsight cluster.

Note: Azure HDInsight is a high-availability service that has redundant gateway nodes, head nodes, and ZooKeeper nodes to keep your HDInsight clusters running smoothly. While this ensures that a single failure will not affect the functionality of a cluster, you may still want to monitor cluster health so you are alerted when an issue does arise. Monitoring cluster health refers to monitoring whether all nodes in your cluster and the components that run on them are available and functioning correctly.

Ambari is the recommended tool for monitoring utilization across the whole cluster. The Ambari dashboard shows easily glanceable widgets that display metrics such as CPU, network, YARN memory, and HDFS disk usage. The specific metrics shown depend on cluster type. The “Hosts” tab shows metrics for individual nodes so you can ensure the load on your cluster is evenly distributed.

References: <https://azure.microsoft.com/en-us/blog/monitoring-on-hdinsight-part-1-an-overview/>

## QUESTION 2

You manage a solution that uses Azure HDInsight clusters.

You need to implement a solution to monitor cluster performance and status.

Which technology should you use?

- A. Azure HDInsight .NET SDK
- B. Azure HDInsight REST API
- C. Ambari REST API

<https://vceplus.com/>

- D. Azure Log Analytics
- E. Ambari Web UI

**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Ambari is the recommended tool for monitoring utilization across the whole cluster. The Ambari dashboard shows easily glanceable widgets that display metrics such as CPU, network, YARN memory, and HDFS disk usage. The specific metrics shown depend on cluster type. The “Hosts” tab shows metrics for individual nodes so you can ensure the load on your cluster is evenly distributed.

The Apache Ambari project is aimed at making Hadoop management simpler by developing software for provisioning, managing, and monitoring Apache Hadoop clusters. Ambari provides an intuitive, easy-to-use Hadoop management web UI backed by its RESTful APIs.

References:

<https://azure.microsoft.com/en-us/blog/monitoring-on-hdinsight-part-1-an-overview/>

<https://ambari.apache.org/>



### QUESTION 3

You configure monitoring for a Microsoft Azure SQL Data Warehouse implementation. The implementation uses PolyBase to load data from comma-separated value (CSV) files stored in Azure Data Lake Gen 2 using an external table.

Files with an invalid schema cause errors to occur.

You need to monitor for an invalid schema error.

For which error should you monitor?

- A. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [com.microsoft.polybase.client.KerberosSecureLogin] occurred while accessing external file.'
- B. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [No FileSystem for scheme: wasbs] occurred while accessing external file.'
- C. Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11": for linked server "(null)", Query aborted- the maximum reject threshold (0 rows) was reached while reading from an external source: 1 rows rejected out of total 1 rows processed.
- D. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge\_Connect: Error [Unable to instantiate LoginClass] occurred while accessing external file.'

<https://vceplus.com/>

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Customer Scenario:

SQL Server 2016 or SQL DW connected to Azure blob storage. The CREATE EXTERNAL TABLE DDL points to a directory (and not a specific file) and the directory contains files with different schemas.

SSMS Error:

Select query on the external table gives the following error:

Msg 7320, Level 16, State 110, Line 14

Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11" for linked server "(null)". Query aborted-- the maximum reject threshold (0 rows) was reached while reading from an external source: 1 rows rejected out of total 1 rows processed.

Possible Reason:

The reason this error happens is because each file has different schema. The PolyBase external table DDL when pointed to a directory recursively reads all the files in that directory. When a column or data type mismatch happens, this error could be seen in SSMS. Possible Solution:

If the data for each table consists of one file, then use the filename in the LOCATION section prepended by the directory of the external files. If there are multiple files per table, put each set of files into different directories in Azure Blob Storage and then you can point LOCATION to the directory instead of a particular file.

The latter suggestion is the best practices recommended by SQLCAT even if you have one file per table.

Incorrect Answers:

A: Possible Reason: Kerberos is not enabled in Hadoop Cluster.

References: <https://techcommunity.microsoft.com/t5/DataCAT/PolyBase-Setup-Errors-and-Possible-Solutions/ba-p/305297>

#### QUESTION 4

**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 questions 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.**

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior.

You need to implement logging.

Solution: Use information stored in Azure Active Directory reports.

Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Instead configure Azure Data Lake Storage diagnostics to store logs and metrics in a storage account.

References: <https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-diagnostic-logs>

#### QUESTION 5

**Note:** This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions 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.**

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior.

You need to implement logging.

Solution: Configure Azure Data Lake Storage diagnostics to store logs and metrics in a storage account.

Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer:** A

**Section:** [none]

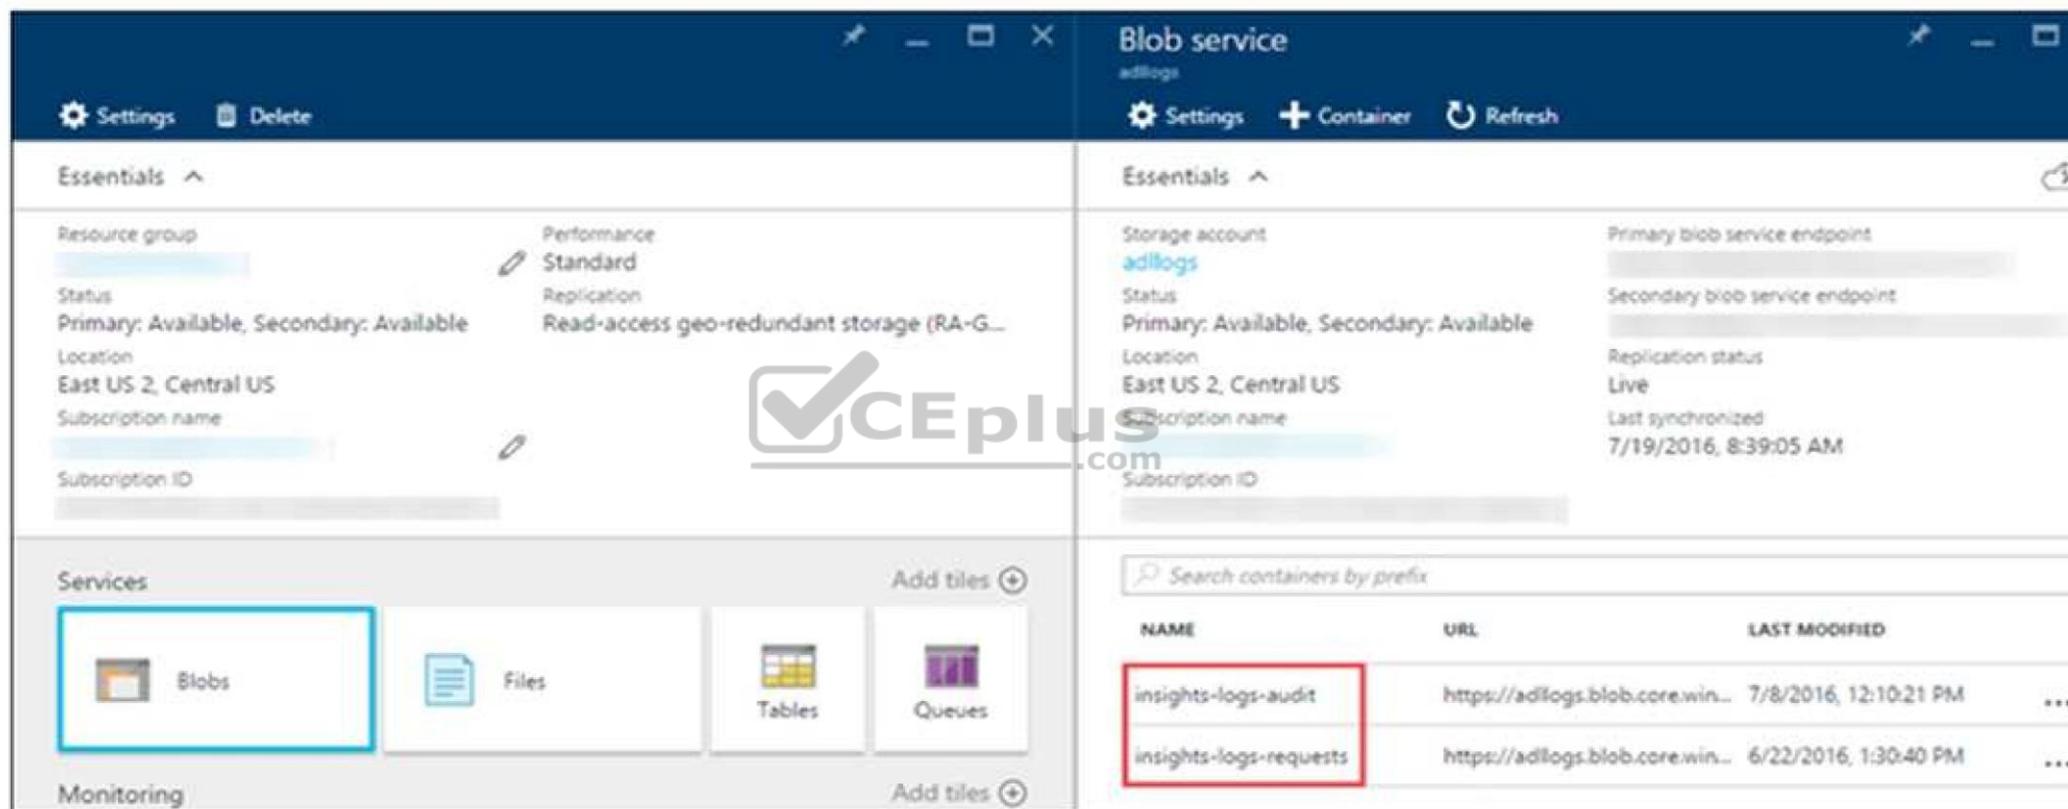
**Explanation**

**Explanation/Reference:**

Explanation:

From the Azure Storage account that contains log data, open the Azure Storage account blade associated with Data Lake Storage Gen1 for logging, and then click Blobs. The Blob service blade lists two containers.

References: <https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-diagnostic-logs>



**QUESTION 6**

**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 questions 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.

A company uses Azure Data Lake Gen 1 Storage to store big data related to consumer behavior.

You need to implement logging.

<https://vceplus.com/>

Solution: Configure an Azure Automation runbook to copy events.

Does the solution meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Instead configure Azure Data Lake Storage diagnostics to store logs and metrics in a storage account.

References: <https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-diagnostic-logs>

#### QUESTION 7

Contoso, Ltd. plans to configure existing applications to use Azure SQL Database.

When security-related operations occur, the security team must be informed.

You need to configure Azure Monitor while minimizing administrative effort.

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

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

- A. Create a new action group to email alerts@contoso.com.
- B. Use alerts@contoso.com as an alert email address.
- C. Use all security operations as a condition.
- D. Use all Azure SQL Database servers as a resource.
- E. Query audit log entries as a condition.

**Correct Answer:** ACD

**Section:** [none]

**Explanation**

**Explanation/Reference:**

<https://vceplus.com/>

References: <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-action-rules>

### QUESTION 8

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a container named Sales in an Azure Cosmos DB database. Sales has 120 GB of data. Each entry in Sales has the following structure.

```
{  
  OrderId: number,  
  OrderDetailId: number,  
  ProductName: string,  
  other information that might vary...  
}
```

The partition key is set to the `OrderId` attribute.

Users report that when they perform queries that retrieve data by `ProductName`, the queries take longer than expected to complete.

You need to reduce the amount of time it takes to execute the problematic queries.

Solution: You create a lookup collection that uses `ProductName` as a partition key.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

<https://vceplus.com/>

Explanation:

One option is to have a lookup collection "ProductName" for the mapping of "ProductName" to "OrderId".

References: <https://azure.microsoft.com/sv-se/blog/azure-cosmos-db-partitioning-design-patterns-part-1/>

#### QUESTION 9

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have a container named Sales in an Azure Cosmos DB database. Sales has 120 GB of data. Each entry in Sales has the following structure.

```
{  
  OrderId: number,  
  OrderDetailId: number,  
  ProductName: string,  
  other information that might vary...  
}
```



The partition key is set to the `OrderId` attribute.

Users report that when they perform queries that retrieve data by `ProductName`, the queries take longer than expected to complete.

You need to reduce the amount of time it takes to execute the problematic queries.

Solution: You create a lookup collection that uses `ProductName` as a partition key and `OrderId` as a value.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

One option is to have a lookup collection "ProductName" for the mapping of "ProductName" to "OrderId".

References: <https://azure.microsoft.com/sv-se/blog/azure-cosmos-db-partitioning-design-patterns-part-1/>

**QUESTION 10**

**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 scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have a container named Sales in an Azure Cosmos DB database. Sales has 120 GB of data. Each entry in Sales has the following structure.

```
{  
  OrderId: number,  
  OrderDetailId: number,  
  ProductName: string,  
  other information that might vary...  
}
```



The partition key is set to the `OrderId` attribute.

Users report that when they perform queries that retrieve data by `ProductName`, the queries take longer than expected to complete.

You need to reduce the amount of time it takes to execute the problematic queries.

Solution: You change the partition key to include `ProductName`.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

One option is to have a lookup collection "ProductName" for the mapping of "ProductName" to "OrderId".

References: <https://azure.microsoft.com/sv-se/blog/azure-cosmos-db-partitioning-design-patterns-part-1/>

**QUESTION 11 HOTSPOT**

You have a new Azure Data Factory environment.

You need to periodically analyze pipeline executions from the last 60 days to identify trends in execution durations. The solution must use Azure Log Analytics to query the data and create charts.

Which diagnostic settings should you configure in Data Factory? To answer, select the appropriate options in the answer area.

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

**Hot Area:**

Answer Area



Log type:  ▼

- ActivityRuns
- AllMetrics
- PipelineRuns
- TriggerRuns

Storage location:  ▼

- An Azure event hub
- An Azure Storage account
- Azure Cosmos DB
- Azure Log Analytics

<https://vceplus.com/>

Correct Answer:

### Answer Area

Log type:

	▼
ActivityRuns	
AllMetrics	
PipelineRuns	
TriggerRuns	

Storage location:

	▼
An Azure event hub	
An Azure Storage account	
Azure Cosmos DB	
Azure Log Analytics	

Section: [none]

#### Explanation

#### Explanation/Reference:

Log type: PipelineRuns

A pipeline run in Azure Data Factory defines an instance of a pipeline execution.

Storage location: An Azure Storage account

Data Factory stores pipeline-run data for only 45 days. Use Monitor if you want to keep that data for a longer time. With Monitor, you can route diagnostic logs for analysis. You can also keep them in a storage account so that you have factory information for your chosen duration.

Save your diagnostic logs to a storage account for auditing or manual inspection. You can use the diagnostic settings to specify the retention time in days.

References: <https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers> <https://docs.microsoft.com/en-us/azure/data-factory/monitor-using-azure-monitor>

<https://vceplus.com/>

### QUESTION 12

You have an Azure SQL database that has masked columns.

You need to identify when a user attempts to infer data from the masked columns.

What should you use?

- A. Azure Advanced Threat Protection (ATP)
- B. custom masking rules
- C. Transparent Data Encryption (TDE)
- D. auditing

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Dynamic Data Masking is designed to simplify application development by limiting data exposure in a set of pre-defined queries used by the application. While Dynamic Data Masking can also be useful to prevent accidental exposure of sensitive data when accessing a production database directly, it is important to note that unprivileged users with ad-hoc query permissions can apply techniques to gain access to the actual data. If there is a need to grant such ad-hoc access, Auditing should be used to monitor all database activity and mitigate this scenario.

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

### QUESTION 13

You have an Azure data solution that contains an Azure SQL data warehouse named DW1.

Several users execute adhoc queries to DW1 concurrently.

You regularly perform automated data loads to DW1.

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

What should you do?

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

<https://vceplus.com/>

D. Assign a smaller resource class to the automated data load queries.

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

To ensure the loading user has enough memory to achieve maximum compression rates, use loading users that are a member of a medium or large resource class.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>

**QUESTION 14**

SIMULATION





Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10277521

You need to generate an email notification to [admin@contoso.com](mailto:admin@contoso.com) if the available storage in an Azure Cosmos DB database named `cosmos10277521` is less than 100,000,000 bytes.

To complete this task, sign in to the Azure portal.

**Correct Answer:** See the explanation below.

<https://vceplus.com/>

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

1. In the Azure portal, click All services, click Azure Cosmos DB, and then click the cosmos10277521 Azure Cosmos DB account.
2. In the resource menu, click Alert Rules to open the Alert rules page.



contoso-2016  
DocumentDB account

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Tutorials and resources

SETTINGS

Replicate data globally

Default consistency

Keys

Properties

Locks

Automation script

MONITORING

Alert rules



<https://vceplus.com/>

3. In the Alert rules page, click Add alert.



4. In the Add an alert rule page, specify:

Metric: Available storage

Condition: Less than

Threshold: 100,000,000 bytes



### Add an alert rule

\* Resource ⓘ  
contoso-2016 (databaseAccounts) ▼

\* Name ⓘ

Description

\* Metric ⓘ  
Available Storage ▼



\* Condition  
greater than ▼

\* Threshold ⓘ  
  
bytes

\* Period ⓘ  
Over the last 30 minutes ▼



Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/monitor-accounts>



## Testlet 2

### Overview

#### Current environment

Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.

The majority of the company's data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

Applications	Tier	Replication	Notes
Internal Contoso	1	Yes	
Internal Contoso	2	SQL Data Sync	Data Sync between databases
Internal Partner	3	Yes	Replicate to Partner
External Contoso	4,5,6	Yes	
External Partner	7,8	No	Partner managed
Internal Distribution and Sales	9	Yes, once ingested at branches	Data ingested from Contoso branches
External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources

The company has a reporting infrastructure that ingests data from local databases and partner services. Partners services consists of distributors, wholesales, and retailers across the world. The company performs daily, weekly, and monthly reporting.

#### Requirements

Tier 3 and Tier 6 through Tier 8 application must use database density on the same server and Elastic pools in a cost-effective manner.

Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.

A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of server going offline.

Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.

- Tier 1 internal applications on the premium P2 tier
- Tier 2 internal applications on the standard S4 tier

The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.

Tier 7 and Tier 8 partner access must be restricted to the database only.

In addition to default Azure backup behavior, Tier 4 and 5 databases must be on a backup strategy that performs a transaction log backup every hour, a differential backup of databases every day and a full backup every week.

Backup strategies must be put in place for all other standalone Azure SQL Databases using Azure SQL-provided backup storage and capabilities.

### Databases

Contoso requires their data estate to be designed and implemented in the Azure Cloud. Moving to the cloud must not inhibit access to or availability of data.

Databases:

Tier 1 Database must implement data masking using the following masking logic:

Data type	Masking requirement
A	Mask 4 or less string data type characters
B	Mask first letter and domain
C	Mask everything except characters at the beginning and end

Tier 2 databases must sync between branches and cloud databases and in the event of conflicts must be set up for conflicts to be won by on-premises databases.

Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.

Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.

A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of a server going offline.

Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.

- Tier 1 internal applications on the premium P2 tier
- Tier 2 internal applications on the standard S4 tier

## Reporting

### Security and monitoring

#### Security

A method of managing multiple databases in the cloud at the same time must be implemented to streamline data management and limit management access to only those requiring access.

#### Monitoring

Monitoring must be set up on every database. Contoso and partners must receive performance reports as part of contractual agreements.

Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.

The Azure SQL Data Warehouse cache must be monitored when the database is being used. A dashboard monitoring key performance indicators (KPIs) indicated by traffic lights must be created and displayed based on the following metrics:

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

Existing Data Protection and Security compliances require that all certificates and keys are internally managed in an on-premises storage.

You identify the following reporting requirements:

- Azure Data Warehouse must be used to gather and query data from multiple internal and external databases
- Azure Data Warehouse must be optimized to use data from a cache
- Reporting data aggregated for external partners must be stored in Azure Storage and be made available during regular business hours in the connecting regions
- Reporting strategies must be improved to real time or near real time reporting cadence to improve competitiveness and the general supply chain
- Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company's main office
- Tier 10 reporting data must be stored in Azure Blobs

#### Issues

Team members identify the following issues:

- Both internal and external client applications run complex joins, equality searches and group-by clauses. Because some systems are managed externally, the queries will not be changed or optimized by Contoso
- External partner organization data formats, types and schemas are controlled by the partner companies

- Internal and external database development staff resources are primarily SQL developers familiar with the Transact-SQL language.
- Size and amount of data has led to applications and reporting solutions not performing are required speeds
- Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
- The company maintains several legacy client applications. Data for these applications remains isolated form other applications. This has led to hundreds of databases being provisioned on a per application basis

### QUESTION 1

**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 questions 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 need to implement diagnostic logging for Data Warehouse monitoring.

Which log should you use?

- A. RequestSteps
- B. DmsWorkers
- C. SqlRequests
- D. ExecRequests



**Correct Answer: C**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Scenario:

The Azure SQL Data Warehouse cache must be monitored when the database is being used.

Metric	Description
A	Low cache hit %, high cache usage %
B	Low cache hit %, low cache usage %
C	High cache hit %, high cache usage %

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-sql-requests-transact-sq>

### QUESTION 2

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

<https://vceplus.com/>

**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 need setup monitoring for tiers 6 through 8.

What should you configure?

- A. extended events for average storage percentage that emails data engineers
- B. an alert rule to monitor CPU percentage in databases that emails data engineers
- C. an alert rule to monitor CPU percentage in elastic pools that emails data engineers
- D. an alert rule to monitor storage percentage in databases that emails data engineers
- E. an alert rule to monitor storage percentage in elastic pools that emails data engineers

**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Scenario:

Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.

Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.

**Testlet 3**

**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 question on 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 on 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

<https://vceplus.com/>

**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

Litware, Inc, is an international car racing and manufacturing company that has 1,000 employees. Most employees are located in Europe. The company supports racing teams that complete in a worldwide racing series.

### Physical Locations

Litware has two main locations: a main office in London, England, and a manufacturing plant in Berlin, Germany.

During each race weekend, 100 engineers set up a remote portable office by using a VPN to connect the datacentre in the London office. The portable office is set up and torn down in approximately 20 different countries each year.

### Existing environment

#### Race Central

During race weekends, Litware uses a primary application named Race Central. Each car has several sensors that send real-time telemetry data to the London datacentre. The data is used for real-time tracking of the cars.

Race Central also sends batch updates to an application named Mechanical Workflow by using Microsoft SQL Server Integration Services (SSIS).

The telemetry data is sent to a MongoDB database. A custom application then moves the data to databases in SQL Server 2017. The telemetry data in MongoDB has more than 500 attributes. The application changes the attribute names when the data is moved to SQL Server 2017.

The database structure contains both OLAP and OLTP databases.

#### Mechanical Workflow

Mechanical Workflow is used to track changes and improvements made to the cars during their lifetime.

Currently, Mechanical Workflow runs on SQL Server 2017 as an OLAP system.

Mechanical Workflow has a named Table1 that is 1 TB. Large aggregations are performed on a single column of Table 1.

### Requirements

### Planned Changes

Litware is the process of rearchitecting its data estate to be hosted in Azure. The company plans to decommission the London datacentre and move all its applications to an Azure datacentre.

### Technical Requirements

Litware identifies the following technical requirements:

- Data collection for Race Central must be moved to Azure Cosmos DB and Azure SQL Database. The data must be written to the Azure datacentre closest to each race and must converge in the least amount of time.
- The query performance of Race Central must be stable, and the administrative time it takes to perform optimizations must be minimized.
- The datacentre for Mechanical Workflow must be moved to Azure SQL data Warehouse.
- Transparent data encryption (TDE) must be enabled on all data stores, whenever possible.
- An Azure Data Factory pipeline must be used to move data from Cosmos DB to SQL Database for Race Central. If the data load takes longer than 20 minutes, configuration changes must be made to Data Factory.
- The telemetry data must migrate toward a solution that is native to Azure.
- The telemetry data must be monitored for performance issues. You must adjust the Cosmos DB Request Units per second (RU/s) to maintain a performance SLA while minimizing the cost of the Ru/s.

### Data Masking Requirements

During rare weekends, visitors will be able to enter the remote portable offices. Litware is concerned that some proprietary information might be exposed. The company identifies the following data masking requirements for the Race Central data that will be stored in SQL Database:

- Only show the last four digits of the values in a column named SuspensionSprings.
- Only Show a zero value for the values in a column named ShockOilWeight.

### QUESTION 1

You are monitoring the Data Factory pipeline that runs from Cosmos DB to SQL Database for Race Central.

You discover that the job takes 45 minutes to run.

What should you do to improve the performance of the job?

- A. Decrease parallelism for the copy activities.
- B. Increase that data integration units.
- C. Configure the copy activities to use staged copy.
- D. Configure the copy activities to perform compression.

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Performance tuning tips and optimization features. In some cases, when you run a copy activity in Azure Data Factory, you see a "Performance tuning tips" message on top of the copy activity monitoring, as shown in the following example. The message tells you the bottleneck that was identified for the given copy run. It also guides you on what to change to boost copy throughput. The performance tuning tips currently provide suggestions like:

- Use PolyBase when you copy data into Azure SQL Data Warehouse.
- Increase Azure Cosmos DB Request Units or Azure SQL Database DTUs (Database Throughput Units) when the resource on the data store side is the bottleneck.
- Remove the unnecessary staged copy.

References: <https://docs.microsoft.com/en-us/azure/data-factory/copy-activity-performance>

## QUESTION 2

What should you implement to optimize SQL Database for Race Central to meet the technical requirements?

- A. the `sp_updatestats` stored procedure
- B. automatic tuning
- C. Query Store
- D. the `dbcc checkdb` command



**Correct Answer: A**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Scenario: The query performance of Race Central must be stable, and the administrative time it takes to perform optimizations must be minimized.

`sp_updatestats` updates query optimization statistics on a table or indexed view. By default, the query optimizer already updates statistics as necessary to improve the query plan; in some cases you can improve query performance by using `UPDATE STATISTICS` or the stored procedure `sp_updatestats` to update statistics more frequently than the default updates.

Incorrect Answers:

D: `dbcc checkdb` checks the logical and physical integrity of all the objects in the specified database

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-updatestats-transact-sql?view=sql-server-ver15>



## Question Set 1

### QUESTION 1

You manage a process that performs analysis of daily web traffic logs on an HDInsight cluster. Each of the 250 web servers generates approximately 10 megabytes (MB) of log data each day. All log data is stored in a single folder in Microsoft Azure Data Lake Storage Gen 2.

You need to improve the performance of the process.

Which two changes should you make? Each correct answer presents a complete solution.

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

- A. Combine the daily log files for all servers into one file
- B. Increase the value of the `mapreduce.map.memory` parameter
- C. Move the log files into folders so that each day's logs are in their own folder
- D. Increase the number of worker nodes
- E. Increase the value of the `hive.tez.container.size` parameter

**Correct Answer:** AC

**Section:** [none]

**Explanation**



#### **Explanation/Reference:**

Explanation:

A: Typically, analytics engines such as HDInsight and Azure Data Lake Analytics have a per-file overhead. If you store your data as many small files, this can negatively affect performance. In general, organize your data into larger sized files for better performance (256MB to 100GB in size). Some engines and applications might have trouble efficiently processing files that are greater than 100GB in size.

C: For Hive workloads, partition pruning of time-series data can help some queries read only a subset of the data which improves performance.

Those pipelines that ingest time-series data, often place their files with a very structured naming for files and folders. Below is a very common example we see for data that is structured by date:

```
\DataSet\YYYY\MM\DD\datafile_YYYY_MM_DD.tsv
```

Notice that the datetime information appears both as folders and in the filename.

References: <https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-performance-tuning-guidance>

### QUESTION 2

You implement an Azure SQL Data Warehouse instance.

<https://vceplus.com/>

You plan to migrate the largest fact table to Azure SQL Data Warehouse. The table resides on Microsoft SQL Server on-premises and is 10 terabytes (TB) in size.

Incoming queries use the primary key Sale Key column to retrieve data as displayed in the following table:

SaleKey	CityKey	CustomerKey	StockItemKey	InvoiceDateKey	Quantity	UnitPrice	TotalExcludingTax
49309	90858	70	69	10/22/13	8	16	128
49313	55710	126	69	10/22/13	2	16	32
49343	44710	234	68	10/22/13	10	16	160
49352	66109	163	70	10/22/13	4	16	64
49448	65312	230	70	10/22/13	8	16	128
49646	85877	271	70	10/24/13	1	16	16
49798	41238	288	69	10/24/13	1	16	16

You need to distribute the large fact table across multiple nodes to optimize performance of the table.

Which technology should you use?

- A. hash distributed table with clustered ColumnStore index
- B. hash distributed table with clustered index
- C. heap table with distribution replicate
- D. round robin distributed table with clustered index
- E. round robin distributed table with clustered ColumnStore index

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Hash-distributed tables improve query performance on large fact tables.

Columnstore indexes can achieve up to 100x better performance on analytics and data warehousing workloads and up to 10x better data compression than traditional rowstore indexes.

Incorrect Answers:

D, E: Round-robin tables are useful for improving loading speed.

<https://vceplus.com/>

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-distribute> <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-query-performance>



<https://vceplus.com/>



<https://vceplus.com/>