



## **Microsoft**

### **Exam Questions DP-200**

Implementing an Azure Data Solution

**NEW QUESTION 1**

- (Exam Topic 1)

You need to ensure that phone-based polling data can be analyzed in the PollingData database.

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

Actions	Answer Area
Parameterize deployment by using Azure Integration Runtime	
Configure an Azure Logic App to deploy the deployment artifact	
Configure Azure DevOps to deploy the deployment artifact	
Create a deployment artifact containing an extracted Azure Resource Manager template	
Parameterize deployment by using the Azure Resource Manager template parameter file	
Create a deployment artifact containing a SQL Server Integration Services (SSIS) package	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

Create a deployment artifact containing an extracted Azure Resource Manager template
Parameterize deployment by using the Azure Resource Manager template parameter file
Configure Azure DevOps to deploy the deployment artifact

Scenario:

All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments  
 No credentials or secrets should be used during deployments

**NEW QUESTION 2**

- (Exam Topic 1)

You need to ensure phone-based polling data upload reliability requirements are met. How should you configure monitoring? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Setting	Value						
Metric	<table border="1"> <tbody> <tr> <td>FileCount</td> <td><input type="checkbox"/></td> </tr> <tr> <td>BlobCapacity</td> <td><input type="checkbox"/></td> </tr> <tr> <td>FileCapacity</td> <td><input type="checkbox"/></td> </tr> </tbody> </table>	FileCount	<input type="checkbox"/>	BlobCapacity	<input type="checkbox"/>	FileCapacity	<input type="checkbox"/>
FileCount	<input type="checkbox"/>						
BlobCapacity	<input type="checkbox"/>						
FileCapacity	<input type="checkbox"/>						
Aggregation	<table border="1"> <tbody> <tr> <td>Avg</td> <td><input type="checkbox"/></td> </tr> <tr> <td>Sum</td> <td><input type="checkbox"/></td> </tr> </tbody> </table>	Avg	<input type="checkbox"/>	Sum	<input type="checkbox"/>		
Avg	<input type="checkbox"/>						
Sum	<input type="checkbox"/>						

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: FileCapacity

FileCapacity is the amount of storage used by the storage account's File service in bytes. Box 2: Avg

The aggregation type of the FileCapacity metric is Avg.

Scenario:

All services and processes must be resilient to a regional Azure outage.

All Azure services must be monitored by using Azure Monitor. On-premises SQL Server performance must be monitored.

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>

### NEW QUESTION 3

- (Exam Topic 2)

You need set up the Azure Data Factory JSON definition for Tier 10 data.

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

NOTE: Each correct selection is worth one point.

Data factory component	Value								
Connector	<table border="1"> <tr><td>connection string</td><td><input type="checkbox"/></td></tr> <tr><td>linked service name string</td><td><input type="checkbox"/></td></tr> <tr><td>gateway connection string</td><td><input type="checkbox"/></td></tr> <tr><td>data store name string</td><td><input type="checkbox"/></td></tr> </table>	connection string	<input type="checkbox"/>	linked service name string	<input type="checkbox"/>	gateway connection string	<input type="checkbox"/>	data store name string	<input type="checkbox"/>
connection string	<input type="checkbox"/>								
linked service name string	<input type="checkbox"/>								
gateway connection string	<input type="checkbox"/>								
data store name string	<input type="checkbox"/>								
Data movement activity	<table border="1"> <tr><td>Azure SQL Data Warehouse</td><td><input type="checkbox"/></td></tr> <tr><td>Azure Files</td><td><input type="checkbox"/></td></tr> <tr><td>Azure Blob</td><td><input type="checkbox"/></td></tr> <tr><td>Azure SQL Database</td><td><input type="checkbox"/></td></tr> </table>	Azure SQL Data Warehouse	<input type="checkbox"/>	Azure Files	<input type="checkbox"/>	Azure Blob	<input type="checkbox"/>	Azure SQL Database	<input type="checkbox"/>
Azure SQL Data Warehouse	<input type="checkbox"/>								
Azure Files	<input type="checkbox"/>								
Azure Blob	<input type="checkbox"/>								
Azure SQL Database	<input type="checkbox"/>								

- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:

Box 1: Connection String

To use storage account key authentication, you use the ConnectionString property, which specifies the information needed to connect to Blob Storage.

Mark this field as a SecureString to store it securely in Data Factory. You can also put account key in Azure Key Vault and pull the accountKey configuration out of the connection string.

Box 2: Azure Blob

Tier 10 reporting data must be stored in Azure Blobs

External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources
---------------------------------	----	---	--

References:

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

### NEW QUESTION 4

- (Exam Topic 2)

You need to set up access to Azure SQL Database for Tier 7 and Tier 8 partners.

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

Actions	Answer Area
Connect to the Database and use Azure PowerShell to create a database firewall rule	
Set the Allow Azure Services to Access Server to Disabled	
In the Azure portal, create a database firewall rule	
In the Azure portal, create a server firewall rule	
Connect to the database and use Transact-SQL to create a database firewall rule	
Set the Allow Azure Services to Access Server setting to Enabled	

- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:

Tier 7 and 8 data access is constrained to single endpoints managed by partners for access Step 1: Set the Allow Azure Services to Access Server setting to Disabled

Set Allow access to Azure services to OFF for the most secure configuration.

By default, access through the SQL Database firewall is enabled for all Azure services, under Allow access to Azure services. Choose OFF to disable access for all Azure services.

Note: The firewall pane has an ON/OFF button that is labeled Allow access to Azure services. The ON setting allows communications from all Azure IP addresses and all Azure subnets. These Azure IPs or subnets might not be owned by you. This ON setting is probably more open than you want your SQL Database to be. The virtual network rule feature offers much finer granular control.

Step 2: In the Azure portal, create a server firewall rule Set up SQL Database server firewall rules

Server-level IP firewall rules apply to all databases within the same SQL Database server. To set up a server-level firewall rule:

➤ In Azure portal, select SQL databases from the left-hand menu, and select your database on the SQL databases page.

➤ On the Overview page, select Set server firewall. The Firewall settings page for the database server opens.

Step 3: Connect to the database and use Transact-SQL to create a database firewall rule

Database-level firewall rules can only be configured using Transact-SQL (T-SQL) statements, and only after you've configured a server-level firewall rule.

To setup a database-level firewall rule:

➤ In Object Explorer, right-click the database and select New Query.

➤ EXECUTE sp\_set\_database\_firewall\_rule N'Example DB Rule','0.0.0.4','0.0.0.4';

➤ On the toolbar, select Execute to create the firewall rule. References:

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

## NEW QUESTION 5

- (Exam Topic 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.

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

**Answer:** E

### 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.

## NEW QUESTION 6

- (Exam Topic 3)

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 efforts

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.

**Answer:** ACE

## NEW QUESTION 7

- (Exam Topic 3)

You are designing a new Lambda architecture on Microsoft Azure. The real-time processing layer must meet the following requirements: Ingestion:

- Receive millions of events per second
- Act as a fully managed Platform-as-a-Service (PaaS) solution
- Integrate with Azure Functions

Stream processing:

- Process on a per-job basis
- Provide seamless connectivity with Azure services
- Use a SQL-based query language

Analytical data store:

- Act as a managed service
- Use a document store
- Provide data encryption at rest

You need to identify the correct technologies to build the Lambda architecture using minimal effort. Which technologies should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Architecture requirement	Answer Area					
<b>Ingestion</b>	<table border="1"> <tr><td>HDInsight Kafka</td><td rowspan="4" style="text-align: center;">v</td></tr> <tr><td>Azure Event Hubs</td></tr> <tr><td>HDInsight Storm</td></tr> <tr><td>HDInsight Spark</td></tr> </table>	HDInsight Kafka	v	Azure Event Hubs	HDInsight Storm	HDInsight Spark
HDInsight Kafka	v					
Azure Event Hubs						
HDInsight Storm						
HDInsight Spark						
<b>Stream Processing</b>	<table border="1"> <tr><td>Azure Stream Analytics</td><td rowspan="4" style="text-align: center;">v</td></tr> <tr><td>HDInsight with Spark Streaming</td></tr> <tr><td>Azure Cosmos DB Change Feed</td></tr> <tr><td>Azure Analysis Services</td></tr> </table>	Azure Stream Analytics	v	HDInsight with Spark Streaming	Azure Cosmos DB Change Feed	Azure Analysis Services
Azure Stream Analytics	v					
HDInsight with Spark Streaming						
Azure Cosmos DB Change Feed						
Azure Analysis Services						
<b>Analytical Data Store</b>	<table border="1"> <tr><td>Hive LLAP on HDInsight</td><td rowspan="4" style="text-align: center;">v</td></tr> <tr><td>Azure Analysis Services</td></tr> <tr><td>Azure Cosmos DB</td></tr> <tr><td>SQL Data Warehouse</td></tr> </table>	Hive LLAP on HDInsight	v	Azure Analysis Services	Azure Cosmos DB	SQL Data Warehouse
Hive LLAP on HDInsight	v					
Azure Analysis Services						
Azure Cosmos DB						
SQL Data Warehouse						

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Azure Event Hubs

This portion of a streaming architecture is often referred to as stream buffering. Options include Azure Event Hubs, Azure IoT Hub, and Kafka.

**NEW QUESTION 8**

- (Exam Topic 3)

Your company manages on-premises Microsoft SQL Server pipelines by using a custom solution.

The data engineering team must implement a process to pull data from SQL Server and migrate it to Azure Blob storage. The process must orchestrate and manage the data lifecycle.

You need to configure Azure Data Factory to connect to the on-premises SQL Server database.

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

Actions	Answer Area
Create an Azure Data Factory resource.	
Configure a self-hosted integration runtime.	
Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.	
Create a database master key on SQL Server.	
Backup the database and send it Azure Blob storage.	
Configure the on-premises SQL Server instance with an integration runtime.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.

You can also use IPsec VPN or Azure ExpressRoute to further secure the communication channel between your on-premises network and Azure.

Azure Virtual Network is a logical representation of your network in the cloud. You can connect an on-premises network to your virtual network by setting up IPsec VPN (site-to-site) or ExpressRoute (private peering).

Step 2: Create an Azure Data Factory resource. Step 3: Configure a self-hosted integration runtime.

You create a self-hosted integration runtime and associate it with an on-premises machine with the SQL Server database. The self-hosted integration runtime is the component that copies data from the SQL Server database on your machine to Azure Blob storage.

Note: 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/tutorial-hybrid-copy-powershell>

**NEW QUESTION 9**

- (Exam Topic 3)

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;

**Answer:** CD

**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>

**NEW QUESTION 10**

- (Exam Topic 3)

A company uses Microsoft Azure SQL Database to store sensitive company data. You encrypt the data and only allow access to specified users from specified locations.

You must monitor data usage, and data copied from the system to prevent data leakage.

You need to configure Azure SQL Database to email a specific user when data leakage occurs.

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

Actions	Answer Area
In Auditing, enable <b>Auditing</b> .	
Configure the service to create alerts for threat detections of type <b>Data Exfiltration</b> .	
In Firewalls and virtual networks, enable <b>Allow access to Azure services</b> .	
Enable advanced threat protection.	
Configure the service to send email alerts to security@contoso.com	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Actions	Answer Area
In Auditing, enable <b>Auditing</b> .	Enable advanced threat protection.
Configure the service to create alerts for threat detections of type <b>Data Exfiltration</b> .	Configure the service to send email alerts to security@contoso.com
In Firewalls and virtual networks, enable <b>Allow access to Azure services</b> .	Configure the service to create alerts for threat detections of type <b>Data Exfiltration</b> .
Enable advanced threat protection.	
Configure the service to send email alerts to security@contoso.com	

**NEW QUESTION 10**

- (Exam Topic 3)

You manage security for a database that supports a line of business application. Private and personal data stored in the database must be protected and encrypted. You need to configure the database to use Transparent Data Encryption (TDE).

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

Actions	Answer Area
Create a database encryption key using a certificate generated with the master key.	
Create a certificate and then create the master key using a password.	
Set the context to the master database.	
Create a master key using a password.	
Set the context to the company database.	
Enable encryption.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a master key  
 Step 2: Create or obtain a certificate protected by the master key  
 Step 3: Set the context to the company database  
 Step 4: Create a database encryption key and protect it by the certificate  
 Step 5: Set the database to use encryption  
 Example code: USE master; GO  
 CREATE MASTER KEY ENCRYPTION BY PASSWORD = '<UseStrongPasswordHere>';  
 go  
 CREATE CERTIFICATE MyServerCert WITH SUBJECT = 'My DEK Certificate'; go  
 USE AdventureWorks2012; GO  
 CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES\_128  
 ENCRYPTION BY SERVER CERTIFICATE MyServerCert; GO  
 ALTER DATABASE AdventureWorks2012 SET ENCRYPTION ON;  
 GO

References:

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

**NEW QUESTION 13**

- (Exam Topic 3)

Your company has on-premises Microsoft SQL Server instance.

The data engineering team plans to implement a process that copies data from the SQL Server instance to Azure Blob storage. The process must orchestrate and manage the data lifecycle.

You need to configure Azure Data Factory to connect to the SQL Server instance.

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

Actions	Answer Area
Configure a linked service to connect to the SQL Server instance.	
From the on-premises network, install and configure a self-hosted integration runtime.	
From the SQL Server, backup the database and then copy the database to Azure Blob storage.	
Deploy an Azure Data Factory	
From the SQL Server, create a database master key	

➤  
➤

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Actions	Answer Area
Configure a linked service to connect to the SQL Server instance.	Deploy an Azure Data Factory.
From the on-premises network, install and configure self-hosted integration runtime.	Configure a linked service to connect to the SQL Server instance.
From the SQL Server, backup the database and then copy the database to Azure Blob storage.	From the on-premises network, install and configure self-hosted integration runtime.
Deploy an Azure Data Factory.	
From the SQL Server, create a database master key	

**NEW QUESTION 14**

- (Exam Topic 3)

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.

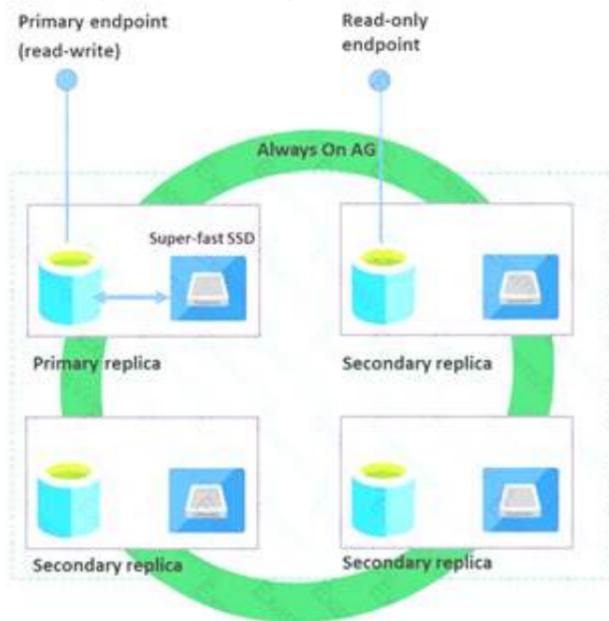
- A. Premium service tier
- B. Virtual machine Scale Sets
- C. Basic service tier
- D. SQL Data Sync
- E. Always On availability groups
- F. Zone-redundant configuration

**Answer:** AEF

**Explanation:**

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.

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

**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>

**NEW QUESTION 18**

- (Exam Topic 3)

You implement an event processing solution using Microsoft Azure Stream Analytics. The solution must meet the following requirements:

- Ingest data from Blob storage
- Analyze data in real time
- Store processed data in Azure Cosmos DB

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

The screenshot shows a question interface with two columns: "Actions" and "Answer Area".

**Actions:**

- Create a query statement with the ORDER BY clause.
- Create a query statement with the SELECT INTO statement
- Configure Blob storage for a reference data JOIN clause
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.
- Set up Cosmos DB as the output
- Configure Blob storage as input, select items with the TIMESTAMP BY clause

**Answer Area:**

Two circular navigation buttons are visible: a right arrow (>) and a left arrow (<).

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

This screenshot is a detailed view of the question interface. A red dashed box highlights the correct answer in the "Answer Area".

**Actions:**

- Create a query statement with the ORDER BY clause
- Create a query statement with the SELECT INTO statement
- Configure Blob storage for a reference data JOIN clause
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.
- Set up Cosmos DB as the output.
- Configure Blob storage as input, select items with the TIMESTAMP BY clause

**Answer Area:**

- Set up Cosmos DB as the output.
- Create a query statement with the SELECT INTO statement
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.

The "Set up Cosmos DB as the output." option is highlighted with a red dashed box, indicating it is the correct answer. Navigation buttons (right arrow > and left arrow <) are also visible.

**NEW QUESTION 19**

- (Exam Topic 3)

You are a data engineer implementing a lambda architecture on Microsoft Azure. You use an open-source big data solution to collect, process, and maintain data. The analytical data store performs poorly.

You must implement a solution that meets the following requirements:

- Provide data warehousing
- Reduce ongoing management activities
- Deliver SQL query responses in less than one second

You need to create an HDInsight cluster to meet the requirements. Which type of cluster should you create?

- A. Interactive Query
- B. Apache Hadoop
- C. Apache HBase
- D. Apache Spark

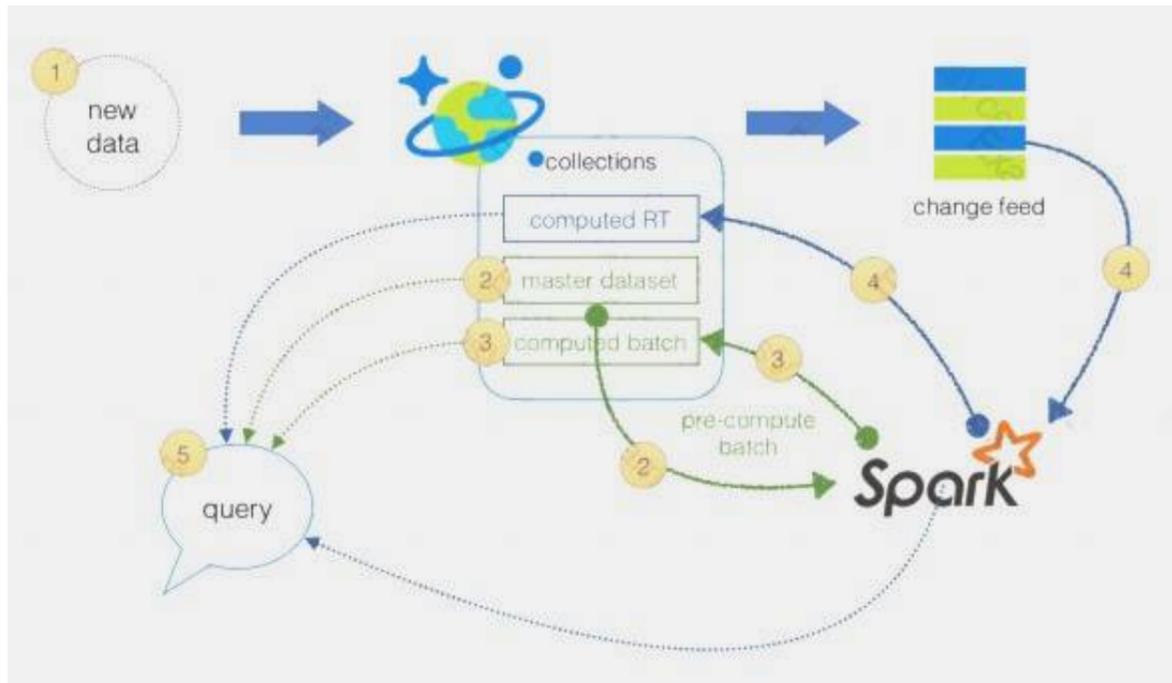
**Answer:** D

**Explanation:**

Lambda Architecture with Azure:

Azure offers you a combination of following technologies to accelerate real-time big data analytics:

- Azure Cosmos DB, a globally distributed and multi-model database service.
- Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.
- The Spark to Azure Cosmos DB Connector



Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.

References:

<https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos->

**NEW QUESTION 24**

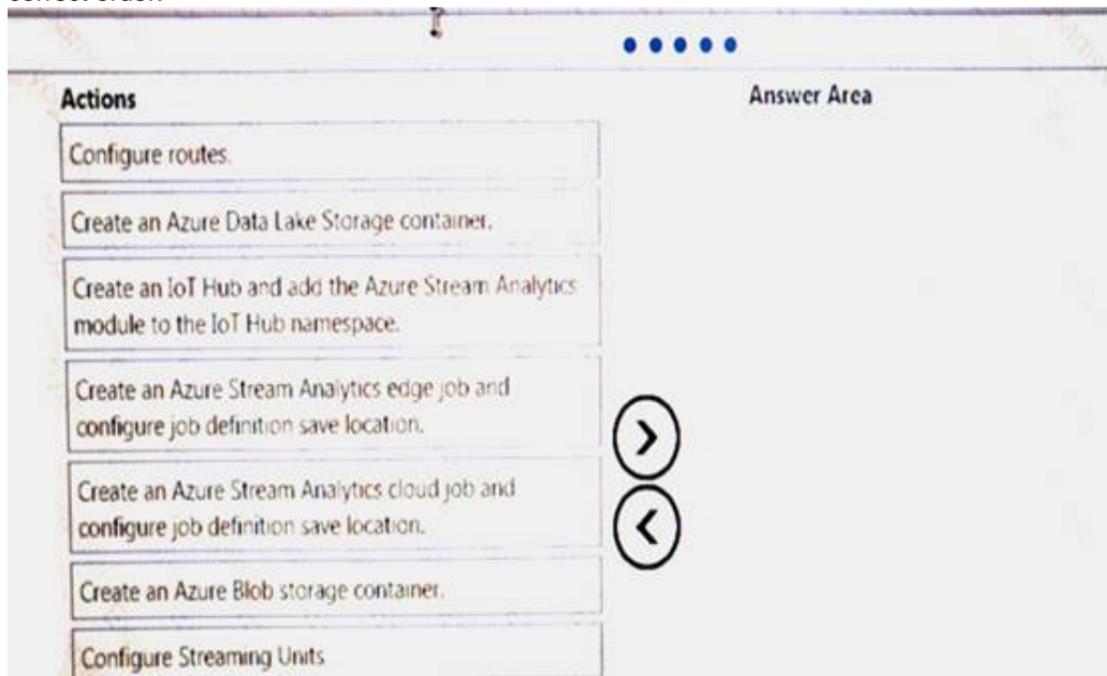
- (Exam Topic 3)

You develop data engineering solutions for a company.

You need to deploy a Microsoft Azure Stream Analytics job for an IoT solution. The solution must:

- Minimize latency.
- Minimize bandwidth usage between the job and IoT device.

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



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

The screenshot shows an exam interface with two columns: 'Actions' and 'Answer Area'. The 'Actions' column contains seven items: 'Configure routes.', 'Create an Azure Data Lake Storage container.', 'Create an IoT Hub and add the Azure Stream Analytics module to the IoT Hub namespace.', 'Create an Azure Stream Analytics edge job and configure job definition save location.', 'Create an Azure Stream Analytics cloud job and configure job definition save location.', 'Create an Azure Blob storage container.', and 'Configure Streaming Units'. The 'Answer Area' contains five items in sequence: 'Configure Streaming Units', 'Create an IoT Hub and add the Azure Stream Analytics module to the IoT Hub namespace.', 'Create an Azure Stream Analytics cloud job and configure job definition save location.', 'Configure routes.', and 'Create an Azure Stream Analytics edge job and configure job definition save location.'. Red dashed boxes and arrows indicate the sequence of actions in the answer area.

**NEW QUESTION 25**

- (Exam Topic 3)

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 Studi
- C. Save the file to an Azure storage account
- D. Export to a BACPAC file by using the Azure portal
- E. Export to a BACPAC file by using Azure PowerShell and save the file locally
- F. Export to a BACPAC file by using the SqlPackage utility

**Answer:** BCE

**NEW QUESTION 30**

- (Exam Topic 3)

You develop data engineering solutions for a company.

A project requires analysis of real-time Twitter feeds. Posts that contain specific keywords must be stored and processed on Microsoft Azure and then displayed by using Microsoft Power BI. You need to implement the solution.

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

The screenshot shows an exam interface with two columns: 'Actions' and 'Answer Area'. The 'Actions' column contains seven items: 'Create an HDInsight cluster with the Hadoop cluster type.', 'Create a Jupyter Notebook.', 'Run a job that uses the Spark Streaming API to ingest data from Twitter.', 'Create a Runbook.', 'Create an HDInsight cluster with the Spark cluster type.', 'Create an table.', and 'Load the hvac table into Power BI Desktop.'. The 'Answer Area' is currently empty.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create an HDInsight cluster with the Spark cluster type Step 2: Create a Jupyter Notebook

Step 3: Create a table

The Jupyter Notebook that you created in the previous step includes code to create an hvac table. Step 4: Run a job that uses the Spark Streaming API to ingest data from Twitter

Step 5: Load the hvac table into Power BI Desktop

You use Power BI to create visualizations, reports, and dashboards from the Spark cluster data. References:

<https://acadgild.com/blog/streaming-twitter-data-using-spark>

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-use-with-data-lake-store>

### NEW QUESTION 32

- (Exam Topic 3)

You develop data engineering solutions for a company. You must migrate data from Microsoft Azure Blob storage to an Azure SQL Data Warehouse for further transformation. You need to implement the solution.

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

Actions	Answer Area
Provision an Azure SQL Data Warehouse instance.	
Connect to the Blob storage container by using SQL Server Management Studio.	
Provision an Azure Blob storage container.	
Run Transact-SQL statements to load data.	
Connect to the Azure SQL Data Warehouse by using SQL Server Management Studio.	
Build external tables by using Azure portal.	
Build external tables by using SQL Server Management Studio.	

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Step 1: Provision an Azure SQL Data Warehouse instance. Create a data warehouse in the Azure portal.

Step 2: Connect to the Azure SQL Data warehouse by using SQL Server Management Studio Connect to the data warehouse with SSMS (SQL Server Management Studio)

Step 3: Build external tables by using the SQL Server Management Studio  
 Create external tables for data in Azure blob storage.

You are ready to begin the process of loading data into your new data warehouse. You use external tables to load data from the Azure storage blob.

Step 4: Run Transact-SQL statements to load data.

You can use the CREATE TABLE AS SELECT (CTAS) T-SQL statement to load the data from Azure Storage Blob into new tables in your data warehouse.

References:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blo>

### NEW QUESTION 36

- (Exam Topic 3)

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 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. Configure 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 AP
- F. Enable multi-region writes.

**Answer:** A

### NEW QUESTION 40

- (Exam Topic 3)

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

**Answer:** B

**NEW QUESTION 41**

- (Exam Topic 3)

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

**Answer:** E

**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>

**NEW QUESTION 44**

- (Exam Topic 3)

Your company uses Microsoft Azure SQL Database configured with Elastic pool. You use Elastic Database jobs to run queries across all databases in the pool. You need to analyze, troubleshoot, and report on components responsible for running Elastic Database jobs. You need to determine the component responsible for running job service tasks.

Which components should you use for each Elastic pool job services task? To answer, drag the appropriate component to the correct task. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**NEW QUESTION 49**

- (Exam Topic 3)

A company manages several on-premises Microsoft SQL Server databases.

You need to migrate the databases to Microsoft Azure by using the backup process of Microsoft SQL Server. Which data technology should you use?

- A. Azure SQL Database Managed Instance
- B. Azure SQL Data Warehouse
- C. Azure Cosmos DB
- D. Azure SQL Database single database

**Answer:** D

**NEW QUESTION 53**

- (Exam Topic 3)

A company builds an application to allow developers to share and compare code. The conversations, code snippets, and links shared by people in the application are stored in a Microsoft Azure SQL Database instance. The application allows for searches of historical conversations and code snippets.

When users share code snippets, the code snippet is compared against previously shared code snippets by using a combination of Transact-SQL functions including SUBSTRING, FIRST\_VALUE, and SQRT. If a match is found, a link to the match is added to the conversation.

Customers report the following issues:

- ▶ Delays occur during live conversations
- ▶ A delay occurs before matching links appear after code snippets are added to conversations

You need to resolve the performance issues.

Which technologies should you use? To answer, drag the appropriate technologies to the correct issues. Each technology may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Technologies	Answer Area	
	Issue	Technology
columnstore index	Delays in conversations	
non-durable table	Delays in match links	
materialized view		
memory-optimized table		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: memory-optimized table

In-Memory OLTP can provide great performance benefits for transaction processing, data ingestion, and transient data scenarios.

Box 2: materialized view

To support efficient querying, a common solution is to generate, in advance, a view that materializes the data in a format suited to the required results set. The Materialized View pattern describes generating prepopulated views of data in environments where the source data isn't in a suitable format for querying, where generating a suitable query is difficult, or where query performance is poor due to the nature of the data or the data store.

These materialized views, which only contain data required by a query, allow applications to quickly obtain the information they need. In addition to joining tables or combining data entities, materialized views can include the current values of calculated columns or data items, the results of combining values or executing transformations on the data items, and values specified as part of the query. A materialized view can even be optimized for just a single query.

References:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/materialized-view>

**NEW QUESTION 57**

- (Exam Topic 3)

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.

Action	Technology					
Extract SQL data on-premises	<table border="1"> <tr><td>Self-hosted integration runtime</td><td rowspan="4" style="text-align: center;">▼</td></tr> <tr><td>Azure-SSIS integration runtime</td></tr> <tr><td>Azure integration runtime</td></tr> <tr><td>Source</td></tr> </table>	Self-hosted integration runtime	▼	Azure-SSIS integration runtime	Azure integration runtime	Source
Self-hosted integration runtime	▼					
Azure-SSIS integration runtime						
Azure integration runtime						
Source						
Load SQL data warehouse	<table border="1"> <tr><td>Self-hosted integration runtime</td><td rowspan="4" style="text-align: center;">▼</td></tr> <tr><td>Azure-SSIS integration runtime</td></tr> <tr><td>Azure integration runtime</td></tr> <tr><td>Sink</td></tr> </table>	Self-hosted integration runtime	▼	Azure-SSIS integration runtime	Azure integration runtime	Sink
Self-hosted integration runtime	▼					
Azure-SSIS integration runtime						
Azure integration runtime						
Sink						

- A. Mastered
- B. Not Mastered

**Answer:** A

**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>

#### NEW QUESTION 60

- (Exam Topic 3)

You manage a process that performs analysis of daily web traffic logs on an HDInsight cluster. Each of 250 web servers generates approximately gigabytes (GB) 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

**Answer:** AC

#### 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>

#### NEW QUESTION 61

- (Exam Topic 3)

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

**Answer:** C

#### 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>

#### NEW QUESTION 64

- (Exam Topic 3)

You manage the Microsoft Azure Databricks environment for a company. You must be able to access a private Azure Blob Storage account. Data must be available to all Azure Databricks workspaces. You need to provide the data access.

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

Actions	Answer Area
Upload a certificate	
Add secrets to the scope	
Use Blob Storage access key	
Create a secret scope	
Configure a JDBC connector	
Mount the Azure Blob Storage container	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a secret scope Step 2: Add secrets to the scope

Note: `dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>")` gets the key that has been stored as a secret in a secret scope.

Step 3: Mount the Azure Blob Storage container

You can mount a Blob Storage container or a folder inside a container through Databricks File System - DBFS. The mount is a pointer to a Blob Storage container, so the data is never synced locally.

Note: To mount a Blob Storage container or a folder inside a container, use the following command:

```
Python dbutils.fs.mount(
source = "wasbs://<your-container-name>@<your-storage-account-name>.blob.core.windows.net", mount_point = "/mnt/<mount-name>",
extra_configs = {"<conf-key>":dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>")}) where:
dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>") gets the key that has been stored as a secret in a secret scope.
```

References:

<https://docs.databricks.com/spark/latest/data-sources/azure/azure-storage.html>

**NEW QUESTION 66**

- (Exam Topic 3)

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.

● ● ● ● ●

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.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**NEW QUESTION 69**

- (Exam Topic 3)

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

You develop data engineering solutions for a company.

A project requires the deployment of resources to Microsoft Azure for batch processing on Azure

HDInsight. Batch processing will run daily and must: Scale to minimize costs

Be monitored for cluster performance

You need to recommend a tool that will monitor clusters and provide information to suggest how to scale. Solution: Download Azure HDInsight cluster logs by using Azure PowerShell.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Reference:

Instead monitor clusters by using Azure Log Analytics and HDInsight cluster management solutions. References: <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-oms-log-analytics-tutorial>

**NEW QUESTION 71**

- (Exam Topic 3)

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.

You need to configure Azure SQL Data Warehouse to receive the data.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**



Answer: A

**NEW QUESTION 86**

- (Exam Topic 3)

You manage a Microsoft Azure SQL Data Warehouse Gen 2.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries

You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

- A. Cache used percentage
- B. Local tempdb percentage
- C. WU percentage
- D. CPU percentage

Answer: B

**NEW QUESTION 91**

- (Exam Topic 3)

You are a data engineer. You are designing a Hadoop Distributed File System (HDFS) architecture. You plan to use Microsoft Azure Data Lake as a data storage repository.

You must provision the repository with a resilient data schema. You need to ensure the resiliency of the Azure Data Lake Storage. What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Requirement	Node				
Provide data access to clients.	<table border="1"> <tr> <td>DataNode</td> <td><input type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input checked="" type="checkbox"/></td> </tr> </table>	DataNode	<input type="checkbox"/>	NameNode	<input checked="" type="checkbox"/>
DataNode	<input type="checkbox"/>				
NameNode	<input checked="" type="checkbox"/>				
Run operations on files and directories of the file system.	<table border="1"> <tr> <td>DataNode</td> <td><input checked="" type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input checked="" type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input checked="" type="checkbox"/>				
NameNode	<input type="checkbox"/>				
Perform block creation, deletion, and replication.	<table border="1"> <tr> <td>DataNode</td> <td><input checked="" type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input checked="" type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input checked="" type="checkbox"/>				
NameNode	<input type="checkbox"/>				

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: NameNode

An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients.

Box 2: DataNode

The DataNodes are responsible for serving read and write requests from the file system's clients. Box 3: DataNode

The DataNodes perform block creation, deletion, and replication upon instruction from the NameNode.

Note: HDFS has a master/slave architecture. An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients. In addition, there are a number of DataNodes, usually one per node in the cluster, which manage storage attached to the nodes that they run on. HDFS exposes a file system namespace and allows user data to be stored in files. Internally, a file is split into one or more blocks and these blocks are stored in a set of DataNodes. The NameNode executes file system namespace operations like opening, closing, and renaming files and directories. It also determines the mapping of blocks to DataNodes. The DataNodes are responsible for serving read and write requests from the file system's clients. The DataNodes also perform block creation, deletion, and replication upon instruction from the NameNode.

References: [https://hadoop.apache.org/docs/r1.2.1/hdfs\\_design.html#NameNode+and+DataNodes](https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html#NameNode+and+DataNodes)

**NEW QUESTION 93**

- (Exam Topic 3)

You manage a financial computation data analysis process. Microsoft Azure virtual machines (VMs) run the process in daily jobs, and store the results in virtual hard drives (VHDs.)

The VMs product results using data from the previous day and store the results in a snapshot of the VHD. When a new month begins, a process creates a new VHD.

You must implement the following data retention requirements:

- Daily results must be kept for 90 days
- Data for the current year must be available for weekly reports
- Data from the previous 10 years must be stored for auditing purposes
- Data required for an audit must be produced within 10 days of a request. You need to enforce the data retention requirements while minimizing cost.

How should you configure the lifecycle policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once, more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segments	Answer Area
<input type="checkbox"/> delete	<pre>{   "version": "0.5",   "rules": [     {       "name": "dataRetention",       "type": "Lifecycle",       "definition": {         "actions": {           "<input type="checkbox"/>": {             "<input type="checkbox"/>": {daysAfterModificationGreaterThan": 365},             "<input type="checkbox"/>": {daysAfterModificationGreaterThan": 3650}           },           "<input type="checkbox"/>": {             "<input type="checkbox"/>": {"daysAfterCreationGreaterThan": 90}           }         }       }     }   ] }</pre>
<input type="checkbox"/> blockBlob	
<input type="checkbox"/> baseBlob	
<input type="checkbox"/> snapshot	
<input type="checkbox"/> tierToCool	
<input type="checkbox"/> tierToArchive	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

The Set-AzStorageAccountManagementPolicy cmdlet creates or modifies the management policy of an Azure Storage account.

Example: Create or update the management policy of a Storage account with ManagementPolicy rule objects.

Action -BaseBlobAction Delete -daysAfterModificationGreaterThan 100

PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToArchive -daysAfterModificationGreaterThan 50

PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToCool -daysAfterModificationGreaterThan 30

PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -SnapshotAction Delete -daysAfterCreationGreaterThan 100

PS C:\>\$filter1 = New-AzStorageAccountManagementPolicyFilter -PrefixMatch ab,cd

PS C:\>\$rule1 = New-AzStorageAccountManagementPolicyRule -Name Test -Action \$action1 -Filter \$filter1

PS C:\>\$action2 = Add-AzStorageAccountManagementPolicyAction -BaseBlobAction Delete

-daysAfterModificationGreaterThan 100

PS C:\>\$filter2 = New-AzStorageAccountManagementPolicyFilter References:

<https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccountmanagementpolicy>

**NEW QUESTION 95**

- (Exam Topic 3)

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

Reports indicate slowdowns during a specific lime 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

**Answer:** B

**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/>

**NEW QUESTION 98**

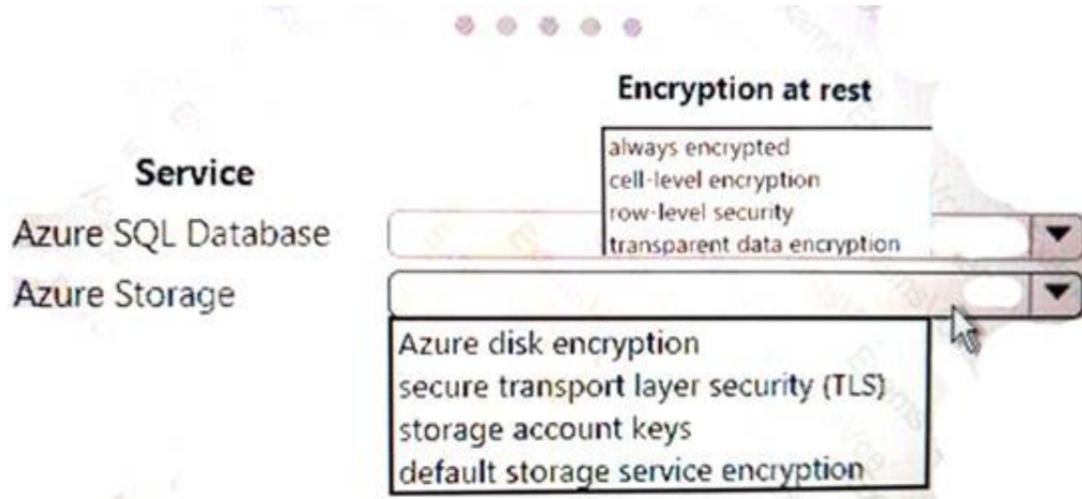
- (Exam Topic 3)

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.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 101**

.....

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

**NEW QUESTION 1**

- (Exam Topic 1)

You need to ensure that phone-based polling data can be analyzed in the PollingData database.

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

Actions	Answer Area
Parameterize deployment by using Azure Integration Runtime	
Configure an Azure Logic App to deploy the deployment artifact	
Configure Azure DevOps to deploy the deployment artifact	
Create a deployment artifact containing an extracted Azure Resource Manager template	
Parameterize deployment by using the Azure Resource Manager template parameter file	
Create a deployment artifact containing a SQL Server Integration Services (SSIS) package	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

Create a deployment artifact containing an extracted Azure Resource Manager template
Parameterize deployment by using the Azure Resource Manager template parameter file
Configure Azure DevOps to deploy the deployment artifact

Scenario:

All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments  
 No credentials or secrets should be used during deployments

**NEW QUESTION 2**

- (Exam Topic 1)

You need to ensure phone-based polling data upload reliability requirements are met. How should you configure monitoring? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Setting	Value						
Metric	<table border="1"> <tbody> <tr> <td>FileCount</td> <td><input type="checkbox"/></td> </tr> <tr> <td>BlobCapacity</td> <td><input type="checkbox"/></td> </tr> <tr> <td>FileCapacity</td> <td><input type="checkbox"/></td> </tr> </tbody> </table>	FileCount	<input type="checkbox"/>	BlobCapacity	<input type="checkbox"/>	FileCapacity	<input type="checkbox"/>
FileCount	<input type="checkbox"/>						
BlobCapacity	<input type="checkbox"/>						
FileCapacity	<input type="checkbox"/>						
Aggregation	<table border="1"> <tbody> <tr> <td>Avg</td> <td><input type="checkbox"/></td> </tr> <tr> <td>Sum</td> <td><input type="checkbox"/></td> </tr> </tbody> </table>	Avg	<input type="checkbox"/>	Sum	<input type="checkbox"/>		
Avg	<input type="checkbox"/>						
Sum	<input type="checkbox"/>						

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: FileCapacity

FileCapacity is the amount of storage used by the storage account's File service in bytes. Box 2: Avg

The aggregation type of the FileCapacity metric is Avg.

Scenario:

All services and processes must be resilient to a regional Azure outage.

All Azure services must be monitored by using Azure Monitor. On-premises SQL Server performance must be monitored.

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>

### NEW QUESTION 3

- (Exam Topic 2)

You need set up the Azure Data Factory JSON definition for Tier 10 data.

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

NOTE: Each correct selection is worth one point.

Data factory component	Value								
Connector	<table border="1"> <tr><td>connection string</td><td><input type="checkbox"/></td></tr> <tr><td>linked service name string</td><td><input type="checkbox"/></td></tr> <tr><td>gateway connection string</td><td><input type="checkbox"/></td></tr> <tr><td>data store name string</td><td><input type="checkbox"/></td></tr> </table>	connection string	<input type="checkbox"/>	linked service name string	<input type="checkbox"/>	gateway connection string	<input type="checkbox"/>	data store name string	<input type="checkbox"/>
connection string	<input type="checkbox"/>								
linked service name string	<input type="checkbox"/>								
gateway connection string	<input type="checkbox"/>								
data store name string	<input type="checkbox"/>								
Data movement activity	<table border="1"> <tr><td>Azure SQL Data Warehouse</td><td><input type="checkbox"/></td></tr> <tr><td>Azure Files</td><td><input type="checkbox"/></td></tr> <tr><td>Azure Blob</td><td><input type="checkbox"/></td></tr> <tr><td>Azure SQL Database</td><td><input type="checkbox"/></td></tr> </table>	Azure SQL Data Warehouse	<input type="checkbox"/>	Azure Files	<input type="checkbox"/>	Azure Blob	<input type="checkbox"/>	Azure SQL Database	<input type="checkbox"/>
Azure SQL Data Warehouse	<input type="checkbox"/>								
Azure Files	<input type="checkbox"/>								
Azure Blob	<input type="checkbox"/>								
Azure SQL Database	<input type="checkbox"/>								

A. Mastered

B. Not Mastered

**Answer: A**

#### Explanation:

Box 1: Connection String

To use storage account key authentication, you use the ConnectionString property, which specifies the information needed to connect to Blob Storage.

Mark this field as a SecureString to store it securely in Data Factory. You can also put account key in Azure Key Vault and pull the accountKey configuration out of the connection string.

Box 2: Azure Blob

Tier 10 reporting data must be stored in Azure Blobs

External Distribution and Sales	10	Yes, once ingested at Contoso main office	Data is ingested from multiple sources
---------------------------------	----	---	--

References:

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

### NEW QUESTION 4

- (Exam Topic 2)

You need to set up access to Azure SQL Database for Tier 7 and Tier 8 partners.

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

Actions	Answer Area
Connect to the Database and use Azure PowerShell to create a database firewall rule	
Set the Allow Azure Services to Access Server to Disabled	
In the Azure portal, create a database firewall rule	
In the Azure portal, create a server firewall rule	
Connect to the database and use Transact-SQL to create a database firewall rule	
Set the Allow Azure Services to Access Server setting to Enabled	

A. Mastered

B. Not Mastered

**Answer: A**

#### Explanation:

Tier 7 and 8 data access is constrained to single endpoints managed by partners for access Step 1: Set the Allow Azure Services to Access Server setting to Disabled

Set Allow access to Azure services to OFF for the most secure configuration.

By default, access through the SQL Database firewall is enabled for all Azure services, under Allow access to Azure services. Choose OFF to disable access for all Azure services.

Note: The firewall pane has an ON/OFF button that is labeled Allow access to Azure services. The ON setting allows communications from all Azure IP addresses and all Azure subnets. These Azure IPs or subnets might not be owned by you. This ON setting is probably more open than you want your SQL Database to be. The virtual network rule feature offers much finer granular control.

Step 2: In the Azure portal, create a server firewall rule Set up SQL Database server firewall rules

Server-level IP firewall rules apply to all databases within the same SQL Database server. To set up a server-level firewall rule:

➤ In Azure portal, select SQL databases from the left-hand menu, and select your database on the SQL databases page.

➤ On the Overview page, select Set server firewall. The Firewall settings page for the database server opens.

Step 3: Connect to the database and use Transact-SQL to create a database firewall rule

Database-level firewall rules can only be configured using Transact-SQL (T-SQL) statements, and only after you've configured a server-level firewall rule.

To setup a database-level firewall rule:

➤ In Object Explorer, right-click the database and select New Query.

➤ EXECUTE sp\_set\_database\_firewall\_rule N'Example DB Rule','0.0.0.4','0.0.0.4';

➤ On the toolbar, select Execute to create the firewall rule. References:

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

## NEW QUESTION 5

- (Exam Topic 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.

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

**Answer:** E

### 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.

## NEW QUESTION 6

- (Exam Topic 3)

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 efforts

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.

**Answer:** ACE

## NEW QUESTION 7

- (Exam Topic 3)

You are designing a new Lambda architecture on Microsoft Azure. The real-time processing layer must meet the following requirements: Ingestion:

- Receive millions of events per second
- Act as a fully managed Platform-as-a-Service (PaaS) solution
- Integrate with Azure Functions

Stream processing:

- Process on a per-job basis
- Provide seamless connectivity with Azure services
- Use a SQL-based query language

Analytical data store:

- Act as a managed service
- Use a document store
- Provide data encryption at rest

You need to identify the correct technologies to build the Lambda architecture using minimal effort. Which technologies should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Architecture requirement	Answer Area					
<b>Ingestion</b>	<table border="1"> <tr><td>HDInsight Kafka</td><td rowspan="4" style="text-align: center;">v</td></tr> <tr><td>Azure Event Hubs</td></tr> <tr><td>HDInsight Storm</td></tr> <tr><td>HDInsight Spark</td></tr> </table>	HDInsight Kafka	v	Azure Event Hubs	HDInsight Storm	HDInsight Spark
HDInsight Kafka	v					
Azure Event Hubs						
HDInsight Storm						
HDInsight Spark						
<b>Stream Processing</b>	<table border="1"> <tr><td>Azure Stream Analytics</td><td rowspan="4" style="text-align: center;">v</td></tr> <tr><td>HDInsight with Spark Streaming</td></tr> <tr><td>Azure Cosmos DB Change Feed</td></tr> <tr><td>Azure Analysis Services</td></tr> </table>	Azure Stream Analytics	v	HDInsight with Spark Streaming	Azure Cosmos DB Change Feed	Azure Analysis Services
Azure Stream Analytics	v					
HDInsight with Spark Streaming						
Azure Cosmos DB Change Feed						
Azure Analysis Services						
<b>Analytical Data Store</b>	<table border="1"> <tr><td>Hive LLAP on HDInsight</td><td rowspan="4" style="text-align: center;">v</td></tr> <tr><td>Azure Analysis Services</td></tr> <tr><td>Azure Cosmos DB</td></tr> <tr><td>SQL Data Warehouse</td></tr> </table>	Hive LLAP on HDInsight	v	Azure Analysis Services	Azure Cosmos DB	SQL Data Warehouse
Hive LLAP on HDInsight	v					
Azure Analysis Services						
Azure Cosmos DB						
SQL Data Warehouse						

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Azure Event Hubs

This portion of a streaming architecture is often referred to as stream buffering. Options include Azure Event Hubs, Azure IoT Hub, and Kafka.

**NEW QUESTION 8**

- (Exam Topic 3)

Your company manages on-premises Microsoft SQL Server pipelines by using a custom solution.

The data engineering team must implement a process to pull data from SQL Server and migrate it to Azure Blob storage. The process must orchestrate and manage the data lifecycle.

You need to configure Azure Data Factory to connect to the on-premises SQL Server database.

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

Actions	Answer Area
Create an Azure Data Factory resource.	
Configure a self-hosted integration runtime.	
Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.	
Create a database master key on SQL Server.	
Backup the database and send it Azure Blob storage.	
Configure the on-premises SQL Server instance with an integration runtime.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a virtual private network (VPN) connection from on-premises to Microsoft Azure.

You can also use IPsec VPN or Azure ExpressRoute to further secure the communication channel between your on-premises network and Azure.

Azure Virtual Network is a logical representation of your network in the cloud. You can connect an on-premises network to your virtual network by setting up IPsec VPN (site-to-site) or ExpressRoute (private peering).

Step 2: Create an Azure Data Factory resource. Step 3: Configure a self-hosted integration runtime.

You create a self-hosted integration runtime and associate it with an on-premises machine with the SQL Server database. The self-hosted integration runtime is the component that copies data from the SQL Server database on your machine to Azure Blob storage.

Note: 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/tutorial-hybrid-copy-powershell>

**NEW QUESTION 9**

- (Exam Topic 3)

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;

**Answer:** CD

**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>

**NEW QUESTION 10**

- (Exam Topic 3)

A company uses Microsoft Azure SQL Database to store sensitive company data. You encrypt the data and only allow access to specified users from specified locations.

You must monitor data usage, and data copied from the system to prevent data leakage.

You need to configure Azure SQL Database to email a specific user when data leakage occurs.

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

Actions	Answer Area
In Auditing, enable <b>Auditing</b> .	
Configure the service to create alerts for threat detections of type <b>Data Exfiltration</b> .	
In Firewalls and virtual networks, enable <b>Allow access to Azure services</b> .	
Enable advanced threat protection.	
Configure the service to send email alerts to security@contoso.com	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Actions	Answer Area
In Auditing, enable <b>Auditing</b> .	Enable advanced threat protection.
Configure the service to create alerts for threat detections of type <b>Data Exfiltration</b> .	Configure the service to send email alerts to security@contoso.com
In Firewalls and virtual networks, enable <b>Allow access to Azure services</b> .	Configure the service to create alerts for threat detections of type <b>Data Exfiltration</b> .
Enable advanced threat protection.	
Configure the service to send email alerts to security@contoso.com	

**NEW QUESTION 10**

- (Exam Topic 3)

You manage security for a database that supports a line of business application. Private and personal data stored in the database must be protected and encrypted. You need to configure the database to use Transparent Data Encryption (TDE).

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

Actions	Answer Area
Create a database encryption key using a certificate generated with the master key.	
Create a certificate and then create the master key using a password.	
Set the context to the master database.	
Create a master key using a password.	
Set the context to the company database.	
Enable encryption.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a master key  
 Step 2: Create or obtain a certificate protected by the master key  
 Step 3: Set the context to the company database  
 Step 4: Create a database encryption key and protect it by the certificate  
 Step 5: Set the database to use encryption  
 Example code: USE master; GO  
 CREATE MASTER KEY ENCRYPTION BY PASSWORD = '<UseStrongPasswordHere>';  
 go  
 CREATE CERTIFICATE MyServerCert WITH SUBJECT = 'My DEK Certificate'; go  
 USE AdventureWorks2012; GO  
 CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES\_128  
 ENCRYPTION BY SERVER CERTIFICATE MyServerCert; GO  
 ALTER DATABASE AdventureWorks2012 SET ENCRYPTION ON;  
 GO

References:

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

**NEW QUESTION 13**

- (Exam Topic 3)

Your company has on-premises Microsoft SQL Server instance.

The data engineering team plans to implement a process that copies data from the SQL Server instance to Azure Blob storage. The process must orchestrate and manage the data lifecycle.

You need to configure Azure Data Factory to connect to the SQL Server instance.

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

Actions	Answer Area
Configure a linked service to connect to the SQL Server instance.	
From the on-premises network, install and configure a self-hosted integration runtime.	
From the SQL Server, backup the database and then copy the database to Azure Blob storage.	
Deploy an Azure Data Factory	
From the SQL Server, create a database master key	

➤  
➤

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Actions	Answer Area
Configure a linked service to connect to the SQL Server instance.	Deploy an Azure Data Factory.
From the on-premises network, install and configure self-hosted integration runtime.	Configure a linked service to connect to the SQL Server instance.
From the SQL Server, backup the database and then copy the database to Azure Blob storage.	From the on-premises network, install and configure self-hosted integration runtime.
Deploy an Azure Data Factory.	
From the SQL Server, create a database master key	

**NEW QUESTION 14**

- (Exam Topic 3)

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.

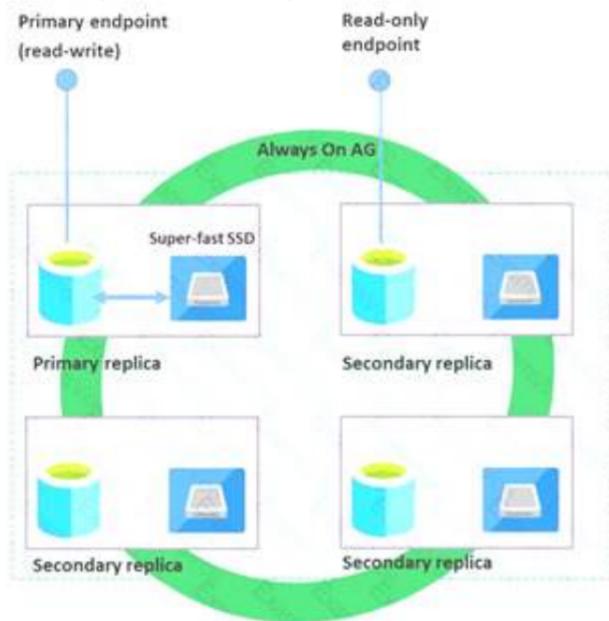
- A. Premium service tier
- B. Virtual machine Scale Sets
- C. Basic service tier
- D. SQL Data Sync
- E. Always On availability groups
- F. Zone-redundant configuration

**Answer:** AEF

**Explanation:**

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.

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

**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>

**NEW QUESTION 18**

- (Exam Topic 3)

You implement an event processing solution using Microsoft Azure Stream Analytics. The solution must meet the following requirements:

- Ingest data from Blob storage
- Analyze data in real time
- Store processed data in Azure Cosmos DB

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

The screenshot shows a question interface with two columns: "Actions" and "Answer Area".

**Actions:**

- Create a query statement with the ORDER BY clause.
- Create a query statement with the SELECT INTO statement
- Configure Blob storage for a reference data JOIN clause
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.
- Set up Cosmos DB as the output
- Configure Blob storage as input, select items with the TIMESTAMP BY clause

**Answer Area:**

Two circular navigation buttons are visible: a right arrow (>) and a left arrow (<).

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

This screenshot is a detailed view of the question interface. A red dashed box highlights the correct answer in the "Answer Area".

**Actions:**

- Create a query statement with the ORDER BY clause
- Create a query statement with the SELECT INTO statement
- Configure Blob storage for a reference data JOIN clause
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.
- Set up Cosmos DB as the output.
- Configure Blob storage as input, select items with the TIMESTAMP BY clause

**Answer Area:**

- Set up Cosmos DB as the output.
- Create a query statement with the SELECT INTO statement
- Configure Azure Event Hub as input; select items with the TIMESTAMP BY clause.

The "Set up Cosmos DB as the output." option is highlighted with a red dashed box, indicating it is the correct answer. Navigation buttons (right arrow > and left arrow <) are also visible.

**NEW QUESTION 19**

- (Exam Topic 3)

You are a data engineer implementing a lambda architecture on Microsoft Azure. You use an open-source big data solution to collect, process, and maintain data. The analytical data store performs poorly.

You must implement a solution that meets the following requirements:

- Provide data warehousing
- Reduce ongoing management activities
- Deliver SQL query responses in less than one second

You need to create an HDInsight cluster to meet the requirements. Which type of cluster should you create?

- A. Interactive Query
- B. Apache Hadoop
- C. Apache HBase
- D. Apache Spark

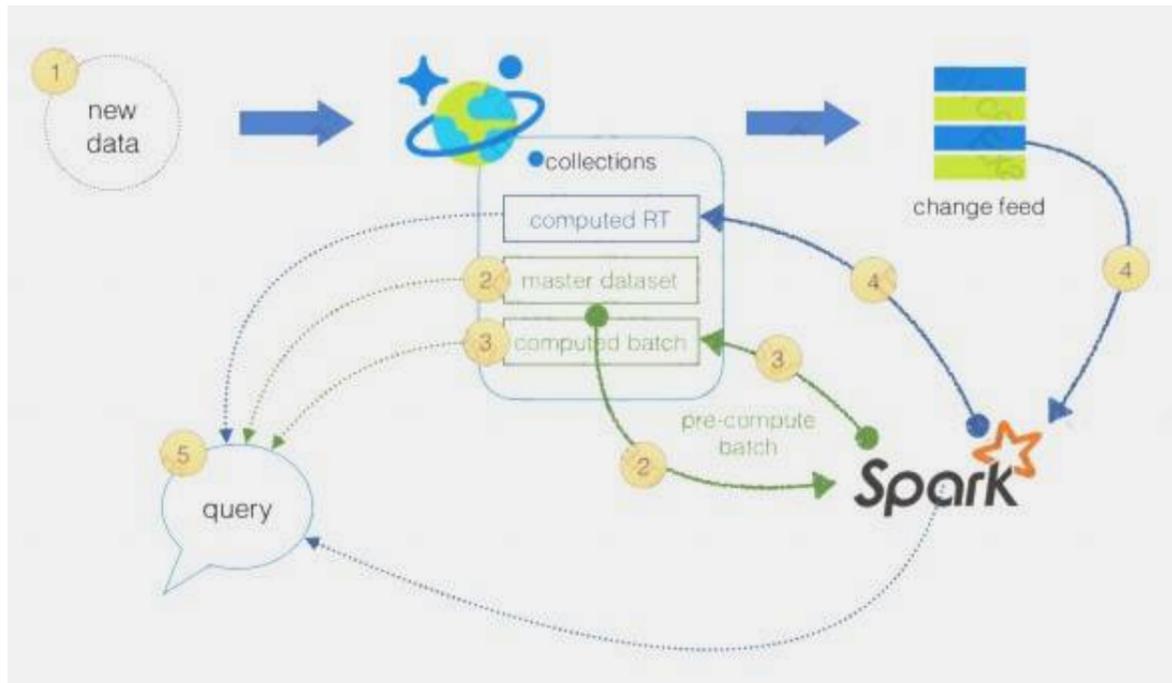
**Answer:** D

**Explanation:**

Lambda Architecture with Azure:

Azure offers you a combination of following technologies to accelerate real-time big data analytics:

- Azure Cosmos DB, a globally distributed and multi-model database service.
- Apache Spark for Azure HDInsight, a processing framework that runs large-scale data analytics applications.
- The Spark to Azure Cosmos DB Connector



Note: Lambda architecture is a data-processing architecture designed to handle massive quantities of data by taking advantage of both batch processing and stream processing methods, and minimizing the latency involved in querying big data.

References:

<https://sqlwithmanoj.com/2018/02/16/what-is-lambda-architecture-and-what-azure-offers-with-its-new-cosmos->

**NEW QUESTION 24**

- (Exam Topic 3)

You develop data engineering solutions for a company.

You need to deploy a Microsoft Azure Stream Analytics job for an IoT solution. The solution must:

- Minimize latency.
- Minimize bandwidth usage between the job and IoT device.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

The screenshot shows an exam interface with two columns: 'Actions' and 'Answer Area'. The 'Actions' column contains seven items: 'Configure routes.', 'Create an Azure Data Lake Storage container.', 'Create an IoT Hub and add the Azure Stream Analytics module to the IoT Hub namespace.', 'Create an Azure Stream Analytics edge job and configure job definition save location.', 'Create an Azure Stream Analytics cloud job and configure job definition save location.', 'Create an Azure Blob storage container.', and 'Configure Streaming Units'. The 'Answer Area' contains five items in sequence: 'Configure Streaming Units', 'Create an IoT Hub and add the Azure Stream Analytics module to the IoT Hub namespace.', 'Create an Azure Stream Analytics cloud job and configure job definition save location.', 'Configure routes.', and 'Create an Azure Stream Analytics edge job and configure job definition save location.'. Red dashed boxes and arrows indicate the sequence of actions in the answer area.

**NEW QUESTION 25**

- (Exam Topic 3)

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 Studi
- C. Save the file to an Azure storage account
- D. Export to a BACPAC file by using the Azure portal
- E. Export to a BACPAC file by using Azure PowerShell and save the file locally
- F. Export to a BACPAC file by using the SqlPackage utility

**Answer:** BCE

**NEW QUESTION 30**

- (Exam Topic 3)

You develop data engineering solutions for a company.

A project requires analysis of real-time Twitter feeds. Posts that contain specific keywords must be stored and processed on Microsoft Azure and then displayed by using Microsoft Power BI. You need to implement the solution.

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

The screenshot shows an exam interface with two columns: 'Actions' and 'Answer Area'. The 'Actions' column contains seven items: 'Create an HDInsight cluster with the Hadoop cluster type.', 'Create a Jupyter Notebook.', 'Run a job that uses the Spark Streaming API to ingest data from Twitter.', 'Create a Runbook.', 'Create an HDInsight cluster with the Spark cluster type.', 'Create an table.', and 'Load the hvac table into Power BI Desktop.'. The 'Answer Area' is currently empty.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create an HDInsight cluster with the Spark cluster type Step 2: Create a Jupyter Notebook

Step 3: Create a table

The Jupyter Notebook that you created in the previous step includes code to create an hvac table. Step 4: Run a job that uses the Spark Streaming API to ingest data from Twitter

Step 5: Load the hvac table into Power BI Desktop

You use Power BI to create visualizations, reports, and dashboards from the Spark cluster data. References:

<https://acadgild.com/blog/streaming-twitter-data-using-spark>

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-use-with-data-lake-store>

### NEW QUESTION 32

- (Exam Topic 3)

You develop data engineering solutions for a company. You must migrate data from Microsoft Azure Blob storage to an Azure SQL Data Warehouse for further transformation. You need to implement the solution.

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

Actions	Answer Area
Provision an Azure SQL Data Warehouse instance.	
Connect to the Blob storage container by using SQL Server Management Studio.	
Provision an Azure Blob storage container.	
Run Transact-SQL statements to load data.	
Connect to the Azure SQL Data Warehouse by using SQL Server Management Studio.	
Build external tables by using Azure portal.	
Build external tables by using SQL Server Management Studio.	

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Step 1: Provision an Azure SQL Data Warehouse instance. Create a data warehouse in the Azure portal.

Step 2: Connect to the Azure SQL Data warehouse by using SQL Server Management Studio Connect to the data warehouse with SSMS (SQL Server Management Studio)

Step 3: Build external tables by using the SQL Server Management Studio  
 Create external tables for data in Azure blob storage.

You are ready to begin the process of loading data into your new data warehouse. You use external tables to load data from the Azure storage blob.

Step 4: Run Transact-SQL statements to load data.

You can use the CREATE TABLE AS SELECT (CTAS) T-SQL statement to load the data from Azure Storage Blob into new tables in your data warehouse.

References:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blo>

### NEW QUESTION 36

- (Exam Topic 3)

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 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. Configure 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 AP
- F. Enable multi-region writes.

**Answer:** A

### NEW QUESTION 40

- (Exam Topic 3)

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

**Answer:** B

**NEW QUESTION 41**

- (Exam Topic 3)

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

**Answer:** E

**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>

**NEW QUESTION 44**

- (Exam Topic 3)

Your company uses Microsoft Azure SQL Database configured with Elastic pool. You use Elastic Database jobs to run queries across all databases in the pod. You need to analyze, troubleshoot, and report on components responsible for running Elastic Database jobs. You need to determine the component responsible for running job service tasks.

Which components should you use for each Elastic pool job services task? To answer, drag the appropriate component to the correct task. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**NEW QUESTION 49**

- (Exam Topic 3)

A company manages several on-premises Microsoft SQL Server databases.

You need to migrate the databases to Microsoft Azure by using the backup process of Microsoft SQL Server. Which data technology should you use?

- A. Azure SQL Database Managed Instance
- B. Azure SQL Data Warehouse
- C. Azure Cosmos DB
- D. Azure SQL Database single database

**Answer:** D

**NEW QUESTION 53**

- (Exam Topic 3)

A company builds an application to allow developers to share and compare code. The conversations, code snippets, and links shared by people in the application are stored in a Microsoft Azure SQL Database instance. The application allows for searches of historical conversations and code snippets.

When users share code snippets, the code snippet is compared against previously shared code snippets by using a combination of Transact-SQL functions including SUBSTRING, FIRST\_VALUE, and SQRT. If a match is found, a link to the match is added to the conversation.

Customers report the following issues:

- Delays occur during live conversations
- A delay occurs before matching links appear after code snippets are added to conversations

You need to resolve the performance issues.

Which technologies should you use? To answer, drag the appropriate technologies to the correct issues. Each technology may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Technologies	Answer Area	
	Issue	Technology
columnstore index	Delays in conversations	
non-durable table	Delays in match links	
materialized view		
memory-optimized table		

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: memory-optimized table

In-Memory OLTP can provide great performance benefits for transaction processing, data ingestion, and transient data scenarios.

Box 2: materialized view

To support efficient querying, a common solution is to generate, in advance, a view that materializes the data in a format suited to the required results set. The Materialized View pattern describes generating prepopulated views of data in environments where the source data isn't in a suitable format for querying, where generating a suitable query is difficult, or where query performance is poor due to the nature of the data or the data store.

These materialized views, which only contain data required by a query, allow applications to quickly obtain the information they need. In addition to joining tables or combining data entities, materialized views can include the current values of calculated columns or data items, the results of combining values or executing transformations on the data items, and values specified as part of the query. A materialized view can even be optimized for just a single query.

References:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/materialized-view>

**NEW QUESTION 57**

- (Exam Topic 3)

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.

Action	Technology					
Extract SQL data on-premises	<table border="1"> <tr><td>Self-hosted integration runtime</td><td rowspan="4">v</td></tr> <tr><td>Azure-SSIS integration runtime</td></tr> <tr><td>Azure integration runtime</td></tr> <tr><td>Source</td></tr> </table>	Self-hosted integration runtime	v	Azure-SSIS integration runtime	Azure integration runtime	Source
Self-hosted integration runtime	v					
Azure-SSIS integration runtime						
Azure integration runtime						
Source						
Load SQL data warehouse	<table border="1"> <tr><td>Self-hosted integration runtime</td><td rowspan="4">v</td></tr> <tr><td>Azure-SSIS integration runtime</td></tr> <tr><td>Azure integration runtime</td></tr> <tr><td>Sink</td></tr> </table>	Self-hosted integration runtime	v	Azure-SSIS integration runtime	Azure integration runtime	Sink
Self-hosted integration runtime	v					
Azure-SSIS integration runtime						
Azure integration runtime						
Sink						

- A. Mastered
- B. Not Mastered

**Answer: A**

**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>

#### NEW QUESTION 60

- (Exam Topic 3)

You manage a process that performs analysis of daily web traffic logs on an HDInsight cluster. Each of 250 web servers generates approximately gigabytes (GB) 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

**Answer:** AC

#### 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>

#### NEW QUESTION 61

- (Exam Topic 3)

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

**Answer:** C

#### 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>

#### NEW QUESTION 64

- (Exam Topic 3)

You manage the Microsoft Azure Databricks environment for a company. You must be able to access a private Azure Blob Storage account. Data must be available to all Azure Databricks workspaces. You need to provide the data access.

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

Actions	Answer Area
Upload a certificate	
Add secrets to the scope	
Use Blob Storage access key	
Create a secret scope	
Configure a JDBC connector	
Mount the Azure Blob Storage container	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a secret scope Step 2: Add secrets to the scope

Note: `dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>")` gets the key that has been stored as a secret in a secret scope.

Step 3: Mount the Azure Blob Storage container

You can mount a Blob Storage container or a folder inside a container through Databricks File System - DBFS. The mount is a pointer to a Blob Storage container, so the data is never synced locally.

Note: To mount a Blob Storage container or a folder inside a container, use the following command:

```
Python dbutils.fs.mount(
source = "wasbs://<your-container-name>@<your-storage-account-name>.blob.core.windows.net", mount_point = "/mnt/<mount-name>",
extra_configs = {"<conf-key>":dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>")}) where:
dbutils.secrets.get(scope = "<scope-name>", key = "<key-name>") gets the key that has been stored as a secret in a secret scope.
```

References:

<https://docs.databricks.com/spark/latest/data-sources/azure/azure-storage.html>

**NEW QUESTION 66**

- (Exam Topic 3)

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.

● ● ● ● ●

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.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**NEW QUESTION 69**

- (Exam Topic 3)

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

You develop data engineering solutions for a company.

A project requires the deployment of resources to Microsoft Azure for batch processing on Azure

HDInsight. Batch processing will run daily and must: Scale to minimize costs

Be monitored for cluster performance

You need to recommend a tool that will monitor clusters and provide information to suggest how to scale. Solution: Download Azure HDInsight cluster logs by using Azure PowerShell.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Reference:

Instead monitor clusters by using Azure Log Analytics and HDInsight cluster management solutions. References: <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-oms-log-analytics-tutorial>

**NEW QUESTION 71**

- (Exam Topic 3)

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.

You need to configure Azure SQL Data Warehouse to receive the data.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**



Answer: A

**NEW QUESTION 86**

- (Exam Topic 3)

You manage a Microsoft Azure SQL Data Warehouse Gen 2.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries

You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

- A. Cache used percentage
- B. Local tempdb percentage
- C. WU percentage
- D. CPU percentage

Answer: B

**NEW QUESTION 91**

- (Exam Topic 3)

You are a data engineer. You are designing a Hadoop Distributed File System (HDFS) architecture. You plan to use Microsoft Azure Data Lake as a data storage repository.

You must provision the repository with a resilient data schema. You need to ensure the resiliency of the Azure Data Lake Storage. What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Requirement	Node				
Provide data access to clients.	<table border="1"> <tr> <td>DataNode</td> <td><input type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input checked="" type="checkbox"/></td> </tr> </table>	DataNode	<input type="checkbox"/>	NameNode	<input checked="" type="checkbox"/>
DataNode	<input type="checkbox"/>				
NameNode	<input checked="" type="checkbox"/>				
Run operations on files and directories of the file system.	<table border="1"> <tr> <td>DataNode</td> <td><input checked="" type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input checked="" type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input checked="" type="checkbox"/>				
NameNode	<input type="checkbox"/>				
Perform block creation, deletion, and replication.	<table border="1"> <tr> <td>DataNode</td> <td><input checked="" type="checkbox"/></td> </tr> <tr> <td>NameNode</td> <td><input type="checkbox"/></td> </tr> </table>	DataNode	<input checked="" type="checkbox"/>	NameNode	<input type="checkbox"/>
DataNode	<input checked="" type="checkbox"/>				
NameNode	<input type="checkbox"/>				

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: NameNode

An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients.

Box 2: DataNode

The DataNodes are responsible for serving read and write requests from the file system's clients. Box 3: DataNode

The DataNodes perform block creation, deletion, and replication upon instruction from the NameNode.

Note: HDFS has a master/slave architecture. An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients. In addition, there are a number of DataNodes, usually one per node in the cluster, which manage storage attached to the nodes that they run on. HDFS exposes a file system namespace and allows user data to be stored in files. Internally, a file is split into one or more blocks and these blocks are stored in a set of DataNodes. The NameNode executes file system namespace operations like opening, closing, and renaming files and directories. It also determines the mapping of blocks to DataNodes. The DataNodes are responsible for serving read and write requests from the file system's clients. The DataNodes also perform block creation, deletion, and replication upon instruction from the NameNode.

References: [https://hadoop.apache.org/docs/r1.2.1/hdfs\\_design.html#NameNode+and+DataNodes](https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html#NameNode+and+DataNodes)

**NEW QUESTION 93**

- (Exam Topic 3)

You manage a financial computation data analysis process. Microsoft Azure virtual machines (VMs) run the process in daily jobs, and store the results in virtual hard drives (VHDs.)

The VMs product results using data from the previous day and store the results in a snapshot of the VHD. When a new month begins, a process creates a new VHD.

You must implement the following data retention requirements:

- Daily results must be kept for 90 days
- Data for the current year must be available for weekly reports
- Data from the previous 10 years must be stored for auditing purposes
- Data required for an audit must be produced within 10 days of a request. You need to enforce the data retention requirements while minimizing cost.

How should you configure the lifecycle policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once, more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segments	Answer Area
<input type="checkbox"/> delete	{
<input type="checkbox"/> blockBob	"version": "0.5",
<input type="checkbox"/> baseBlob	"rules": [
<input type="checkbox"/> snapshot	{
<input type="checkbox"/> tierToCool	"name": "dataRetention",
<input type="checkbox"/> tierToArchive	"type": "Lifecycle",
	"definition": {
	"actions": {
	"": {
	"": { "daysAfterModificationGreaterThan": 365 },
	"": { "daysAfterModificationGreaterThan": 3650 }
	},
	"": {
	"": { "daysAfterCreationGreaterThan": 90 }
	}
	}
	}
	}

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

The Set-AzStorageAccountManagementPolicy cmdlet creates or modifies the management policy of an Azure Storage account.

Example: Create or update the management policy of a Storage account with ManagementPolicy rule objects.

Action -BaseBlobAction Delete -daysAfterModificationGreaterThan 100

PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToArchive -daysAfterModificationGreaterThan 50

PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -BaseBlobAction TierToCool -daysAfterModificationGreaterThan 30

PS C:\>\$action1 = Add-AzStorageAccountManagementPolicyAction -InputObject \$action1 -SnapshotAction Delete -daysAfterCreationGreaterThan 100

PS C:\>\$filter1 = New-AzStorageAccountManagementPolicyFilter -PrefixMatch ab,cd

PS C:\>\$rule1 = New-AzStorageAccountManagementPolicyRule -Name Test -Action \$action1 -Filter \$filter1

PS C:\>\$action2 = Add-AzStorageAccountManagementPolicyAction -BaseBlobAction Delete

-daysAfterModificationGreaterThan 100

PS C:\>\$filter2 = New-AzStorageAccountManagementPolicyFilter References:

<https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccountmanagementpolicy>

**NEW QUESTION 95**

- (Exam Topic 3)

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

Reports indicate slowdowns during a specific lime 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

**Answer:** B

**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/>

**NEW QUESTION 98**

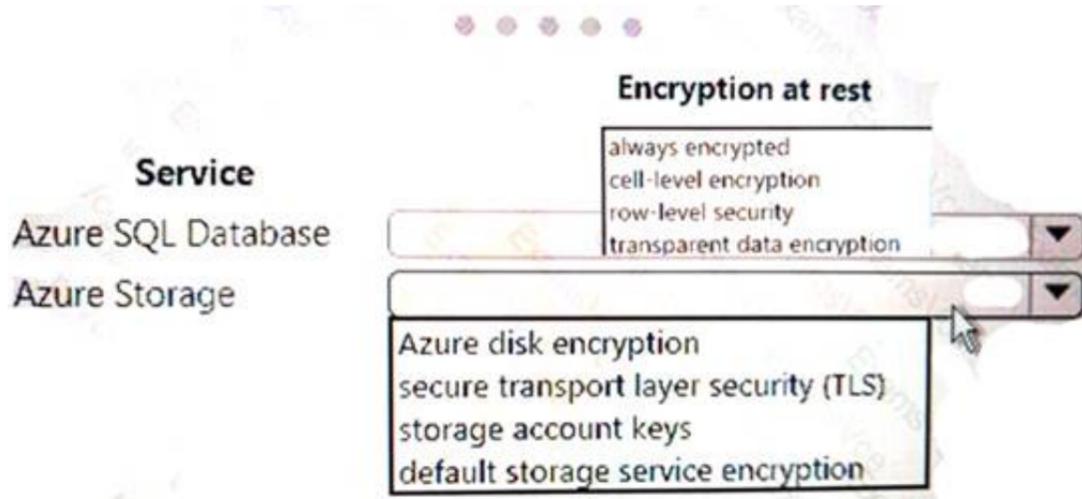
- (Exam Topic 3)

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.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 101**

.....

## Relate Links

**100% Pass Your DP-200 Exam with ExamBible Prep Materials**

<https://www.exambible.com/DP-200-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>