



# DP-200<sup>Q&As</sup>

Implementing an Azure Data Solution

## Pass Microsoft DP-200 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.lead4pass.com/dp-200.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft  
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers



### QUESTION 1

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

A single storage account to store all operations including reads, writes and deletes

Retention of an on-premises copy of historical operations

You need to configure the storage account.

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AB

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

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

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

### QUESTION 2

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

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

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.


- A. Premium service tier
- B. Virtual machine Scale Sets

- C. Basic service tier
- D. SQL Data Sync
- E. Always On availability groups
- F. Zone-redundant configuration

Correct Answer: AEF

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

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



## Business Critical service tier: collocated compute and storage

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

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

### QUESTION 3

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: BE

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

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

References: <https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction>  
<https://www.mongodb.com/scale/types-of-nosql-databases>

---

#### QUESTION 4

A company is deploying a service-based data environment. You are developing a solution to process this data.

The solution must meet the following requirements:

Use an Azure HDInsight cluster for data ingestion from a relational database in a different cloud service

Use an Azure Data Lake Storage account to store processed data

Allow users to download processed data

You need to recommend technologies for the solution.

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

Hot Area:



## Answer Area

### Data process

### Technology

Ingest

RevoScaleR	✓
Apache Sqoop	
Apache DistCp	
Azure CLI	

Process

Apache DistCp	✓
Apache Kafka	
C#	
Apache Hive	

Download

Apache Sqoop	✓
MapReduce	
RevoScaleR	
Ambari Hive View	

Correct Answer:



## Answer Area

Data process	Technology								
Ingest	<table><tr><td>RevoScaleR</td><td>✓</td></tr><tr><td>Apache Sqoop</td><td></td></tr><tr><td>Apache DistCp</td><td></td></tr><tr><td>Azure CLI</td><td></td></tr></table>	RevoScaleR	✓	Apache Sqoop		Apache DistCp		Azure CLI	
RevoScaleR	✓								
Apache Sqoop									
Apache DistCp									
Azure CLI									
Process	<table><tr><td>Apache DistCp</td><td>✓</td></tr><tr><td>Apache Kafka</td><td></td></tr><tr><td>C#</td><td></td></tr><tr><td>Apache Hive</td><td></td></tr></table>	Apache DistCp	✓	Apache Kafka		C#		Apache Hive	
Apache DistCp	✓								
Apache Kafka									
C#									
Apache Hive									
Download	<table><tr><td>Apache Sqoop</td><td>✓</td></tr><tr><td>MapReduce</td><td></td></tr><tr><td>RevoScaleR</td><td></td></tr><tr><td>Ambari Hive View</td><td></td></tr></table>	Apache Sqoop	✓	MapReduce		RevoScaleR		Ambari Hive View	
Apache Sqoop	✓								
MapReduce									
RevoScaleR									
Ambari Hive View									

### Box 1: Apache Sqoop

Apache Sqoop is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured datastores such as relational databases.

Azure HDInsight is a cloud distribution of the Hadoop components from the Hortonworks Data Platform (HDP).

Incorrect Answers:

DistCp (distributed copy) is a tool used for large inter/intra-cluster copying. It uses MapReduce to effect its distribution, error handling and recovery, and reporting. It expands a list of files and directories into input to map tasks, each of which

will copy a partition of the files specified in the source list. Its MapReduce pedigree has endowed it with some quirks in both its semantics and execution.

RevoScaleR is a collection of proprietary functions in Machine Learning Server used for practicing data science at scale. For data scientists, RevoScaleR gives you data-related functions for import, transformation and manipulation,

summarization, visualization, and analysis.

### Box 2: Apache Kafka

Apache Kafka is a distributed streaming platform.

A streaming platform has three key capabilities:



Publish and subscribe to streams of records, similar to a message queue or enterprise messaging system.

Store streams of records in a fault-tolerant durable way.

Process streams of records as they occur.

Kafka is generally used for two broad classes of applications:

Building real-time streaming data pipelines that reliably get data between systems or applications

Building real-time streaming applications that transform or react to the streams of data

Box 3: Ambari Hive View

You can run Hive queries by using Apache Ambari Hive View. The Hive View allows you to author, optimize, and run Hive queries from your web browser.

References:

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

<https://kafka.apache.org/intro>

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-use-hive-ambari-view>

---

## QUESTION 5

A company plans to analyze a continuous flow of data from a social media platform by using Microsoft Azure Stream Analytics. The incoming data is formatted as one record per row.

You need to create the input stream.

How should you complete the REST API segment? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:





```
{
  "properties": {
    "type": "stream",
    "serialization": {
      1
    },
    "properties": {
      "fieldDelimiter": ".",
      "encoding": "UTF8"
    }
  },
  "datasource": {
    2
    "properties": {
      "serviceBusNamespace": "sampleServiceBus",
      "shareAccessPolicyName": "SampleReceiver",
      "shareAccessPolicyKey": "<policykey>"
    },
    "eventHubName": "sampleEventHub"
  }
}
```

1 ▼
"type": "CSV"
"type": "Avro",
"type": "JSON",

2 ▼
"type": "Microsoft.Storage/Blob",
"type": "Microsoft.ServiceBus/EventHub",
"type": "Microsoft.Devices/IoTHubs",

Correct Answer:



```
{
  "properties": {
    "type": "stream",
    "serialization": {
      1
    },
    "properties": {
      "fieldDelimiter": ",",
      "encoding": "UTF8"
    }
  },
  "datasource": {
    2
    "properties": {
      "serviceBusNamespace": "sampleServiceBus",
      "shareAccessPolicyName": "SampleReceiver",
      "shareAccessPolicyKey": "<policykey>"
      "eventHubName": "sampleEventHub"
    }
  }
}
```

1 ▼
"type": "CSV"
"type": "Avro",
"type": "JSON",

2 ▼
"type": "Microsoft.Storage/Blob",
"type": "Microsoft.ServiceBus/EventHub",
"type": "Microsoft.Devices/IoTHubs",

#### QUESTION 6

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

You need to migrate the databases to Microsoft Azure by using a backup process of Microsoft SQL Server.



Which data technology should you use?

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

Correct Answer: D

Managed instance is a new deployment option of Azure SQL Database, providing near 100% compatibility with the latest SQL Server on-premises (Enterprise Edition) Database Engine, providing a native virtual network (VNet) implementation that addresses common security concerns, and a business model favorable for on-premises SQL Server customers. The managed instance deployment model allows existing SQL Server customers to lift and shift their on-premises applications to the cloud with minimal application and database changes.

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

#### QUESTION 7

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.

Hot Area:

#### Answer Area

#### Requirement

#### Node

Provide data access to clients.

DataNode	V
NameNode	

Run operations on files and directories of the file system.

DataNode	V
NameNode	

Perform block creation, deletion, and replication.

DataNode	V
NameNode	



Correct Answer:

### Answer Area

#### Requirement

#### Node

Provide data access to clients.

DataNode	✓
NameNode	

Run operations on files and directories of the file system.

DataNode	✓
NameNode	

Perform block creation, deletion, and replication.

DataNode	✓
NameNode	

Explanation/Reference:

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

### QUESTION 8

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

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

Which tool should you use?

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



E. Data Migration Assistant (DMA)

Correct Answer: E

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

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

---

## QUESTION 9

You plan to deploy an Azure Cosmos DB database that supports multi-master replication.

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

Provide a recovery point objective (RPO) of less than 15 minutes.

Provide a recovery time objective (RTO) of zero minutes.

What are three possible consistency levels that you can select? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Correct Answer: CDE





Region(s)	Replication mode	Consistency level	RPO	RTO
1	Single or Multi-Master	Any Consistency Level	< 240 Minutes	< 1 Week
>1	Single Master	Session, Consistent Prefix, Eventual	< 15 minutes	< 15 minutes
>1	Single Master	Bounded Staleness	K & T	< 15 minutes
>1	Single Master	Strong	0	< 15 minutes
>1	Multi-Master	Session, Consistent Prefix, Eventual	< 15 minutes	0
>1	Multi-Master	Bounded Staleness	K & T	0

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

#### QUESTION 10

You manage a solution that uses Azure HDInsight clusters.

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

Which technology should you use?

- A. Azure HDInsight .NET SDK
- B. Azure HDInsight REST API
- C. Ambari REST API
- D. Azure Log Analytics
- E. Ambari Web UI

Correct Answer: E

Ambari is the recommended tool for monitoring utilization across the whole cluster. The Ambari dashboard shows easily glanceable widgets that display metrics such as CPU, network, YARN memory, and HDFS disk usage. The specific



metrics shown depend on cluster type. The "Hosts" tab shows metrics for individual nodes so you can ensure the load on your cluster is evenly distributed.

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

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

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

## QUESTION 11

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

The solution has the following requirements:

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

You need to provision storage for the solution.

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

Select and Place:

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

Correct Answer:





<b>Enable encryption on the Azure Data Lake using the Azure portal.</b>	<b>Select and configure an encryption key storage container</b>
	<b>Add an access policy for the new Azure Data Lake account to the key storage container</b>
	<b>Create a new Azure Data Lake Storage account with Azure Data Lake managed encryption keys</b>
	<b>Create a new Azure Data Lake Storage account with encryption disabled.</b>
<b>Create a new Azure Data Lake Storage account with Azure Key Vault managed encryption keys</b>	

## QUESTION 12

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

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

You need to configure data encryption for external applications.

Solution:

1.

Access the Always Encrypted Wizard in SQL Server Management Studio

2.

Select the column to be encrypted

3.

Set the encryption type to Randomized

4.

Configure the master key to use the Windows Certificate Store

5.

Validate configuration results and deploy the solution

Does the solution meet the goal?



A. Yes

B. No

Correct Answer: B

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

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

The screenshot shows the 'Master Key Configuration' window. On the left is a navigation pane with links: Introduction, Column Selection, Master Key Configuration (selected), Validation, Summary, and Results. The main area contains instructions: 'To generate a new column encryption key, a column master key must be selected to protect it. The column master key is stored outside of the database.' Below this, there are two sections. The first, 'Select column master key:', has a dropdown menu set to 'Auto generate column master key'. The second, 'Select the key store provider:', has two radio buttons: 'Windows certificate store' (unselected) and 'Azure Key Vault' (selected). Below the radio buttons, it says 'You are signed in as sstein@microsoft.com.' with a 'Change user' link. Then, 'Select a subscription to use:' followed by a dropdown menu. Finally, 'Select an Azure Key Vault:' followed by a dropdown menu showing 'AeKeyVault'. This last dropdown is highlighted with a red rectangle. At the bottom right are buttons for '< Previous', 'Next >', and 'Cancel'.

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

### QUESTION 13

You need to process and query ingested Tier 9 data.



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

NOTE: Each correct selection is worth one point.

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

Correct Answer: EF

Explanation:

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

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

1.

Create a Kafka enabled Event Hubs namespace.

2.

Create a Kafka client that sends messages to the event hub.

3.

Create a Stream Analytics job that copies data from the event hub into an Azure blob storage. Scenario:

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

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

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

[Latest DP-200 Dumps](#)

[DP-200 VCE Dumps](#)

[DP-200 Study Guide](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

## Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.  
You can view Vendor list of All Certification Exams offered:

<https://www.lead4pass.com/allproducts>

## Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <b>One Year Free Update</b> Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.	 <b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.	 <b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © lead4pass, All Rights Reserved.