



DP-201^{Q&As}

Designing an Azure Data Solution

Pass Microsoft DP-201 Exam with 100% Guarantee

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

<https://www.lead4pass.com/dp-201.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

You have an Azure Storage account.

You plan to copy one million image files to the storage account.

You plan to share the files with an external partner organization. The partner organization will analyze the files during the next year.

You need to recommend an external access solution for the storage account. The solution must meet the following requirements:

Ensure that only the partner organization can access the storage account.

Ensure that access of the partner organization is removed automatically after 365 days.

What should you include in the recommendation?

- A. shared keys
- B. Azure Blob storage lifecycle management policies
- C. Azure policies
- D. shared access signature (SAS)

Correct Answer: D

A shared access signature (SAS) is a URI that grants restricted access rights to Azure Storage resources. You can provide a shared access signature to clients who should not be trusted with your storage account key but to whom you wish to delegate access to certain storage account resources. By distributing a shared access signature URI to these clients, you can grant them access to a resource for a specified period of time, with a specified set of permissions.

Reference: <https://docs.microsoft.com/en-us/rest/api/storageservices/delegate-access-with-shared-access-signature>

QUESTION 2

DRAG DROP

You are designing an Azure SQL Data Warehouse for a financial services company. Azure Active Directory will be used to authenticate the users.

You need to ensure that the following security requirements are met:

1.

Department managers must be able to create new database.

2.

The IT department must assign users to databases.

3.



Permissions granted must be minimized.

Which role memberships should you recommend? To answer, drag the appropriate roles to the correct groups. Each role may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view

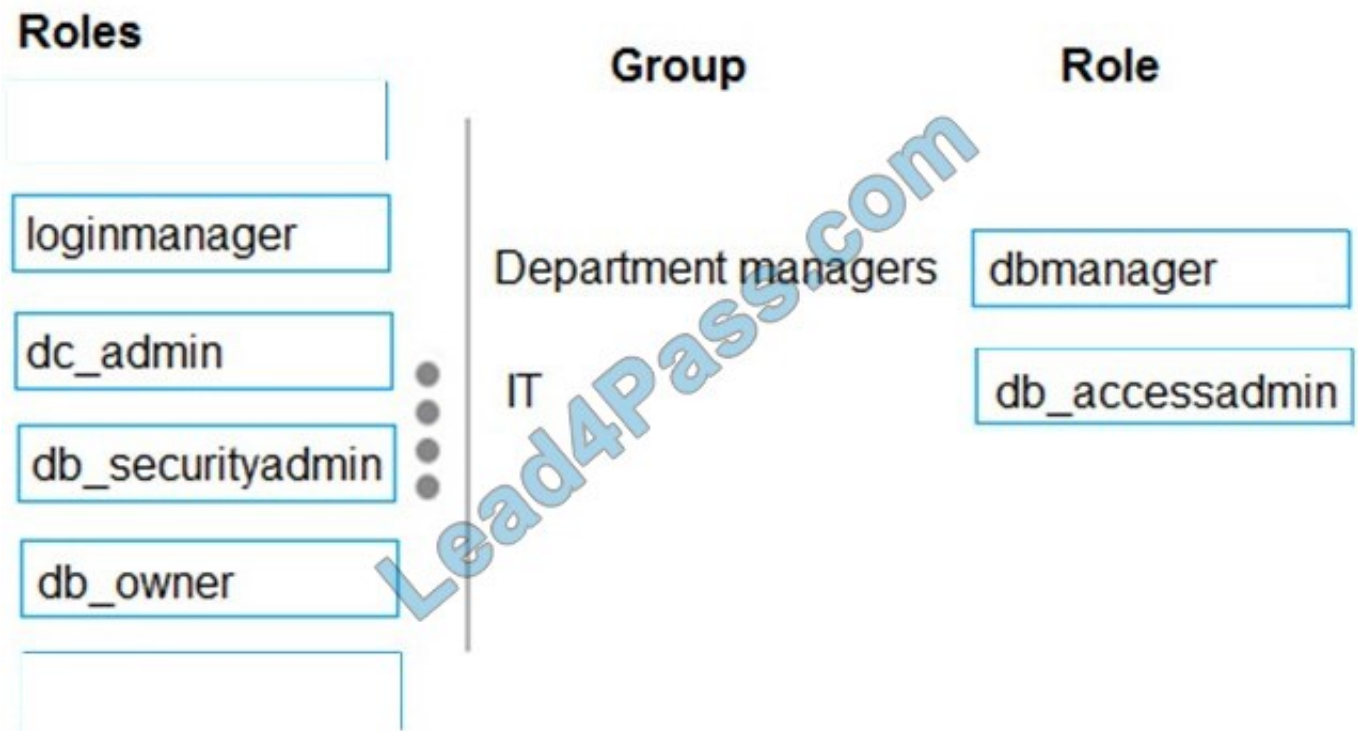
content.

NOTE: Each correct selection is worth one point.

Select and Place:

| Roles | Group | Role |
|------------------|---------------------------|------|
| dbmanager | Department managers IT | |
| loginmanager | | |
| dc_admin | | |
| db_securityadmin | | |
| db_owner | | |
| db_accessadmin | | |

Correct Answer:



Box 1: dbmanager

Members of the dbmanager role can create new databases.

Box 2: db_accessadmin

Members of the db_accessadmin fixed database role can add or remove access to the database for Windows logins, Windows groups, and SQL Server logins.

References:

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

QUESTION 3

You need to optimize storage for CONT_SQL3.

What should you recommend?

- A. AlwaysOn
- B. Transactional processing
- C. General
- D. Data warehousing

Correct Answer: B

CONT_SQL3 with the SQL Server role, 100 GB database size, Hyper-VM to be migrated to Azure VM. The storage



should be configured to optimized storage for database OLTP workloads.

Azure SQL Database provides three basic in-memory based capabilities (built into the underlying database engine) that can contribute in a meaningful way to performance improvements:

In-Memory Online Transactional Processing (OLTP) Clustered columnstore indexes intended primarily for Online Analytical Processing (OLAP) workloads Nonclustered columnstore indexes geared towards Hybrid Transactional/Analytical Processing (HTAP) workloads

References: <https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sql-database.html>

QUESTION 4

You need to design network access to the SQL Server data.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

| Solution component | Value |
|---------------------------------------|---|
| Tool | <input type="text" value=""/> 1433 8080 |
| SQL Server Network Configuration port | <input type="text" value=""/> SQL Server Configuration Manager SQL Server Management Studio |

Correct Answer:



Answer Area

| Solution component | Value |
|---------------------------------------|--|
| Tool | <input type="text" value="1433"/> <input type="text" value="8080"/> |
| SQL Server Network Configuration port | <input type="text" value="SQL Server Configuration Manager"/> <input type="text" value="SQL Server Management Studio"/> |

Box 1: 8080

1433 is the default port, but we must change it as CONT_SQL3 must not communicate over the default ports. Because port 1433 is the known standard for SQL Server, some organizations specify that the SQL Server port number should be

changed to enhance security.

Box 2: SQL Server Configuration Manager

You can configure an instance of the SQL Server Database Engine to listen on a specific fixed port by using the SQL Server Configuration Manager.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-a-server-to-listen-on-a-specific-tcp-port?view=sql-server-2017>

QUESTION 5

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

You have an Azure SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.



You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You create a SELECT trigger on the table in SQL Database that writes the query to a new table in the database, and then executes a stored procedure that looks up the column classifications and joins to the query text.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Instead add classifications to the columns that contain sensitive data and turn on Auditing.

Note: Auditing has been enhanced to log sensitivity classifications or labels of the actual data that were returned by the query. This would enable you to gain insights on who is accessing sensitive data.

References:

<https://azure.microsoft.com/en-us/blog/announcing-public-preview-of-data-discovery-classification-for-microsoft-azure-sql-data-warehouse/>

QUESTION 6

You are designing an Azure Cosmos DB database that will support vertices and edges. Which Cosmos DB API should you include in the design?

- A. SQL
- B. Cassandra
- C. Gremlin
- D. Table

Correct Answer: C

The Azure Cosmos DB Gremlin API can be used to store massive graphs with billions of vertices and edges.

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

QUESTION 7

You need to recommend a storage solution for a sales system that will receive thousands of small files per minute. The files will be in JSON, text, and CSV formats. The files will be processed and transformed before they are loaded into an Azure data warehouse. The files must be stored and secured in folders.

Which storage solution should you recommend?

- A. Azure Data Lake Storage Gen2

- B. Azure Cosmos DB
- C. Azure SQL Database
- D. Azure Blob storage

Correct Answer: A

Azure provides several solutions for working with CSV and JSON files, depending on your needs. The primary landing place for these files is either Azure Storage or Azure Data Lake Store.¹

Azure Data Lake Storage is an optimized storage for big data analytics workloads.

Incorrect Answers:

D: Azure Blob Storage containers is a general purpose object store for a wide variety of storage scenarios. Blobs are stored in containers, which are similar to folders.

References: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/scenarios/csv-and-json>

QUESTION 8

A company stores sensitive information about customers and employees in Azure SQL Database.

You need to ensure that the sensitive data remains encrypted in transit and at rest.

What should you recommend?

- A. Transparent Data Encryption
- B. Always Encrypted with secure enclaves
- C. Azure Disk Encryption
- D. SQL Server AlwaysOn

Correct Answer: B

Incorrect Answers:

A: Transparent Data Encryption (TDE) encrypts SQL Server, Azure SQL Database, and Azure SQL Data Warehouse data files, known as encrypting data at rest. TDE does not provide encryption across communication channels.

References: <https://cloudblogs.microsoft.com/sqlserver/2018/12/17/confidential-computing-using-always-encrypted-with-secure-enclaves-in-sql-server-2019-preview/>

QUESTION 9

A company has a real-time data analysis solution that is hosted on Microsoft Azure. The solution uses Azure Event Hub to ingest data and an Azure Stream Analytics cloud job to analyze the data. The cloud job is configured to use 120

Streaming Units (SU).



You need to optimize performance for the Azure Stream Analytics job.

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

NOTE: Each correct selection is worth one point.

- A. Implement event ordering
- B. Scale the SU count for the job up
- C. Implement Azure Stream Analytics user-defined functions (UDF)
- D. Scale the SU count for the job down
- E. Implement query parallelization by partitioning the data output
- F. Implement query parallelization by partitioning the data output

Correct Answer: BF

Scale out the query by allowing the system to process each input partition separately.

F: A Stream Analytics job definition includes inputs, a query, and output. Inputs are where the job reads the data stream from.

Reference: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-parallelization>

QUESTION 10

You are designing a big data storage solution. The solution must meet the following requirements:

1.
Provide unlimited account sizes.
2.
Support a hierarchical file system.
3.
Be optimized for parallel analytics workloads. Which storage solution should you use?

- A. Azure Data Lake Storage Gen2
- B. Azure Blob storage
- C. Apache HBase in Azure HDInsight
- D. Azure Cosmos DB

Correct Answer: A

Azure Data Lake Storage is optimized performance for parallel analytics workloads



A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized.

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

QUESTION 11

HOTSPOT

You are planning the deployment of two separate Azure Cosmos DB databases named db1 and db2.

You need to recommend a deployment strategy that meets the following requirements:

1. Costs for both databases must be minimized.
2. Db1 must meet an SLA of 99.99% for both reads and writes.
3. Db2 must meet an SLA of 99.99% for writes and 99.999% for reads.

Which deployment strategy should you recommend for each database? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Db1:

| | |
|--|---|
| | ▼ |
| A single read/write region | |
| A single write region and multiple read regions | |
| Multiple read regions and multiple write regions | |

Db2:

| | |
|--|---|
| | ▼ |
| A single read/write region | |
| A single write region and multiple read regions | |
| Multiple read regions and multiple write regions | |

Correct Answer:



Answer Area

Db1: ▼

- A single read/write region
- A single write region and multiple read regions
- Multiple read regions and multiple write regions

Db2: ▼

- A single read/write region
- A single write region and multiple read regions
- Multiple read regions and multiple write regions

Db1: A single read/write region Db2: A single write region and multi read regions

| Operation type | Single region | Multi-region (single region writes) | Multi-region (multi-region writes) |
|----------------|---------------|-------------------------------------|------------------------------------|
| Writes | 99.99 | 99.99 | 99.999 |
| Reads | 99.99 | 99.999 | 99.999 |

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

QUESTION 12

You need to ensure that emergency road response vehicles are dispatched automatically.

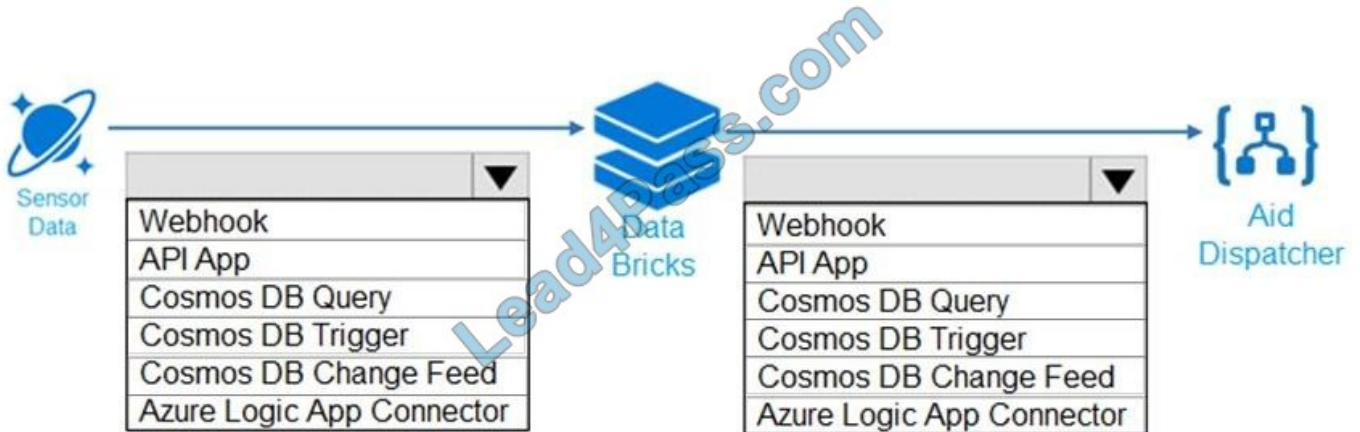
How should you design the processing system? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

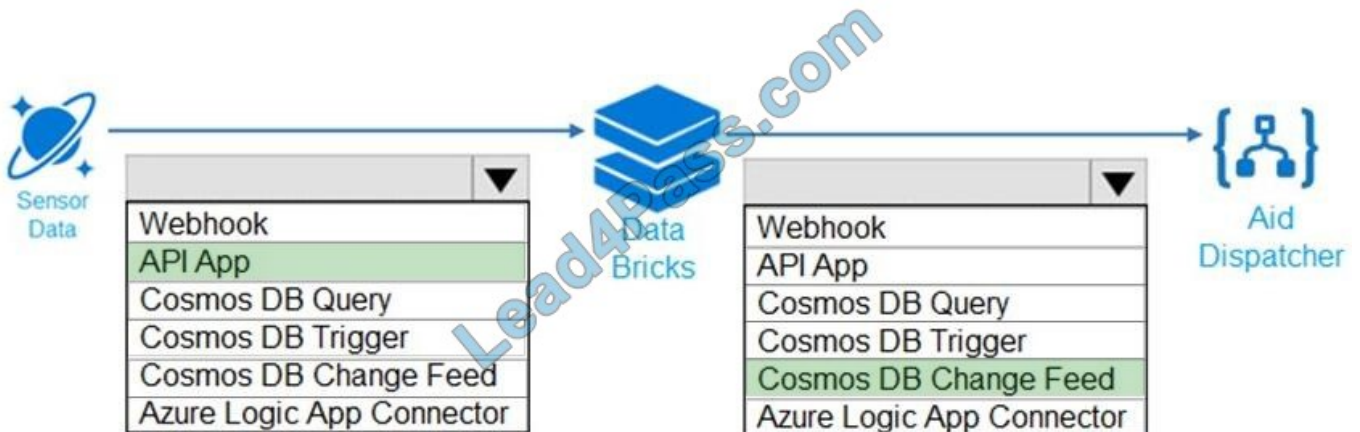


Answer Area




Correct Answer:

Answer Area



Explanation: Box1: API App



Events generated from the IoT data sources are sent to the stream ingestion layer through Azure HDInsight Kafka as a stream of messages. HDInsight Kafka stores streams of data in topics for a configurable of time.

Kafka consumer, Azure Databricks, picks up the message in real time from the Kafka topic, to process the data based on the business logic and can then send to Serving layer for storage.

Downstream storage services, like Azure Cosmos DB, Azure SQL Data warehouse, or Azure SQL DB, will then be a



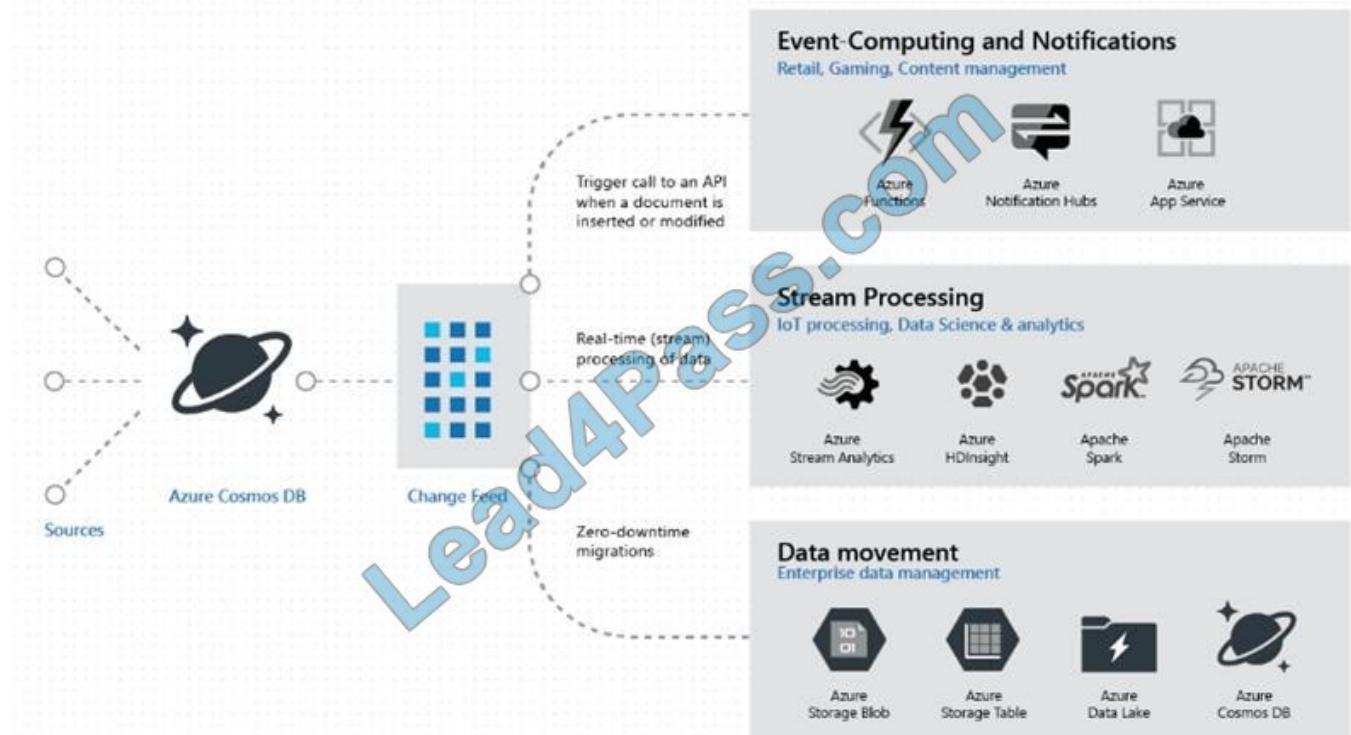
data source for presentation and action layer.

Business analysts can use Microsoft Power BI to analyze warehoused data. Other applications can be built upon the serving layer as well. For example, we can expose APIs based on the service layer data for third party uses.

Box 2: Cosmos DB Change Feed

Change feed support in Azure Cosmos DB works by listening to an Azure Cosmos DB container for any changes. It then outputs the sorted list of documents that were changed in the order in which they were modified.

The change feed in Azure Cosmos DB enables you to build efficient and scalable solutions for each of these patterns, as shown in the following image:



QUESTION 13

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

You have streaming data that is received by Azure Event Hubs and stored in Azure Blob storage. The data contains social media posts that relate to a keyword of Contoso.

You need to count how many times the Contoso keyword and a keyword of Litware appear in the same post every 30 seconds. The data must be available to Microsoft Power BI in near real-time.

Solution: You create an Azure Stream Analytics job that uses an input from Event Hubs to count the posts that have the specified keywords, then and send the data to an Azure SQL database. You consume the data in Power BI by using

DirectQuery mode.



Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

Reference: <https://docs.microsoft.com/en-us/power-bi/service-real-time-streaming>

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-twitter-sentiment-analysis-trends>

[DP-201 PDF Dumps](#)

[DP-201 VCE Dumps](#)

[DP-201 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:



| | | |
|---|---|--|
|  <p>One Year Free Update 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.</p> |  <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p> |  <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p> |
|---|---|--|

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.