Microsoft

Exam Questions DP-900

Microsoft Azure Data Fundamentals

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
Select the answer that correctly completes the sentence.

Select the answer that correctly completes the sentence.

Answer Area

| Microsoft Power BI Desktop ▼ | provides a full range of Power BI data modeling and report editing features. |

Microsoft Power BI Desktop
The Microsoft Power BI service
The Microsoft Power BI phone app

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text, letter Description automatically generated

NEW QUESTION 2
To complete the sentence, select the appropriate option in the answer area.

A visualization that illustrates a university's current student enrollment per department is an example of _____ ▼ analytics.

cognitive
descriptive
predictive
prescriptive

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text, table Description automatically generated
Generally speaking, data analytics comes in four types:
» Descriptive, to answer the question: What's happening?
» Diagnostic, to answer the question: Why's happening?
» Predictive, to answer the question: What will happen?
» Prescriptive, to answer the question: What actions should we take? Company name Description automatically generated

NEW QUESTION 3
You have the following SQL query.

```
INSERT INTO dbo.Products (ProductID, ProductName, Price, ProductDescription)
    VALUES (1, 'Clamp', 12.48, 'Workbench clamp') ;
```

NOTE: Each correct selection is worth one point.

Answer Area

Dbo.Products:	A column
	A database
	A table
	An index

ProductName:	A column
	A database
	A table
	An index

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Dbo.Products:	A column
	A database
	A table
	An index

ProductName:	**A column**
	A database
	A table
	An index

NEW QUESTION 4
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

In batch processing, | delay in delivering data processing results is acceptable. | ▼

NEW QUESTION 5
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Azure Databricks is an Apache Spark-based collaborative analytics platform.	○	○
Azure Analysis Services is used for transactional workloads.	○	○
Azure Data Factory orchestrates data integration workflows.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 6
Select the answer that correctly completes the sentence.

Answer Area

Column family databases ▼	natively support the analysis of relationships between entities.
Column family databases	
Document databases	
Graph databases	
Key-value stores	

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Column family databases ▼	natively support the analysis of relationships between entities.
Column family databases	
Document databases	
Graph databases	
Key-value stores	

NEW QUESTION 7
Which Azure Data Factory component provides the compute environment for activities?

A. a linked service
B. an integration runtime
C. a control flow
D. a pipeline

Answer: B

Explanation:
The Integration Runtime (IR) is the compute infrastructure used by Azure Data Factory to provide the following data integration capabilities across different network environments:
≫ Data Flow
≫ Data movement
≫ Activity dispatch
≫ SSIS package execution

NEW QUESTION 8
To complete the sentence, select the appropriate option in the answer area.

Answer Area

A visualization that shows a university's current student enrollment
versus the maximum capacity is an example of

cognitive	analytics.
descriptive	
predictive	
prescriptive	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Generally speaking, data analytics comes in four types (Figure 1): Descriptive, to answer the question: What's happening? Diagnostic, to answer the question: Why's happening? Predictive, to answer the question: What will happen?
Prescriptive, to answer the question: What actions should we take?

NEW QUESTION 9
Which three actions are used when performing diagnostic analytics?

A. • Identify data anomalies• Collect anomaly-related data• Discover relationships and trends that explain anomalies.
B. • Identify data anomalies• Predict when anomalies are likely to recur• Use analytical tools such as decision trees and regression.
C. • Apply serf-learning loops to draw inferences from anomalies.• Predict when anomalies are likely to recur• Use analytical tools such as decision trees and regression.
D. • Identify data anomalies• Use machine learning to find patterns in large datasets.• Discover relationships and trends that explain anomalies.

Answer: A

NEW QUESTION 10
What are two uses of data visualization? Each correct answer present a complete solution. NOTE: Each correct selection is worth one point.

A. Represent trends and patterns over time.
B. Communicate the significance of data.
C. implement machine learning to predict future values.
D. Consistently implement business logic across reports

Answer: AB

Explanation:
Data visualization is a key component in being able to gain insight into your data. It helps make big and small data easier for humans to understand. It also makes it easier to detect patterns, trends, and outliers in groups of data.
Data visualization brings data to help you find key business insights quickly and effectively.

NEW QUESTION 10
Which scenario is an example of a streaming workload?

A. sending transactions that are older than a month to an archive
B. sending transactions daily from point of sale (POS) devices
C. sending telemetry data from edge devices
D. sending cloud infrastructure metadata every 30 minutes

Answer: C

NEW QUESTION 14
Which statement is an example of Data Manipulation Language (DML)?

A. INSERT
B. DISABLE
C. ALTER
D. DROP

Answer: A

Explanation:
Data Manipulation Language (DML) statements:
➢ DELETE
➢ INSERT
➢ UPDATE

NEW QUESTION 15
To complete the sentence, select the appropriate option in the answer area.

Relational data is stored in [▼]

a file system as unstructured data.
a hierarchal folder structure.
a tabular form of rows and columns.
comma-separated value (CSV) files.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated

NEW QUESTION 20
To complete the sentence, select the appropriate option in the answer area.

Transparent Data Encryption (TDE) encrypts ▼

a column to protect data at rest and in transit.
queries and their results in order to protect data in transit.
the database to protect data at rest.
the server to protect data at rest.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Transparent Data Encryption (TDE) encrypts ▼

a column to protect data at rest and in transit.
queries and their results in order to protect data in transit.
the database to protect data at rest.
the server to protect data at rest.

NEW QUESTION 25
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
Azure Table storage supports multiple read replicas.	○	○
Azure Table storage supports multiple write regions.	○	○
The Azure Cosmos DB Table API supports multiple read replicas.	○	○
The Azure Cosmos DB Table API supports multiple write regions.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
Azure Table storage supports multiple read replicas.	○	○
Azure Table storage supports multiple write regions.	○	○
The Azure Cosmos DB Table API supports multiple read replicas.	○	○
The Azure Cosmos DB Table API supports multiple write regions.	○	○

NEW QUESTION 28
To complete the sentence, select the appropriate option in the answer area.

Descriptive analytics tells you ▼

| what is most likely to occur in the future. |
| what occurred in the past. |
| which actions you can perform to affect outcomes. |
| why something occurred in the past. |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Descriptive analytics tells you ▼

| what is most likely to occur in the future. |
| what occurred in the past. |
| which actions you can perform to affect outcomes. |
| why something occurred in the past. |

NEW QUESTION 29
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Batch processing can output data to a file store	○	○
Batch processing can output data to a relational database	○	○
Batch processing can output data to a NoSQL database	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Big data solutions often use long-running batch jobs to filter, aggregate, and otherwise prepare the data for analysis. Usually these jobs involve reading source files from scalable storage (like HDFS, Azure Data Lake Store, and Azure Storage), processing them, and writing the output to new files in scalable storage.
Box 2: No
Box 3: No

NEW QUESTION 34
Match the types of data stores to the appropriate scenarios.
To answer, drag the appropriate data store type from the column on the left to its scenario on the right. Each data store type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Data Store Types	Answer Area	
Graph		Application users and their default language
Object		Medical images and their associated metadata
Key/value		Employee data that shows the relationships between employees

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Data Store Types	Answer Area	
Graph	Key/value	Application users and their default language
Object	Object	Medical images and their associated metadata
Key/value	Graph	Employee data that shows the relationships between employees

NEW QUESTION 39
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	○	○
Partition keys are used in Azure Cosmos DB to optimize queries.	○	○
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
The API determines the type of account to create. Azure Cosmos DB provides five APIs: Core (SQL) and MongoDB for document data, Gremlin for graph data, Azure Table, and Cassandra. Currently, you must create a separate account for each API.
Box 2: Yes
Azure Cosmos DB uses partitioning to scale individual containers in a database to meet the performance needs of your application. In partitioning, the items in a container are divided into distinct subsets called logical partitions. Logical partitions are formed based on the value of a partition key that is associated with each item in a container.
Box 3: No
Logical partitions are formed based on the value of a partition key that is associated with each item in a container.

NEW QUESTION 44
Your company needs to ensure that in Azure virtual machine can connect to Azure SQL databases without exposing the databases to the internet. What should you use?

A. Azure DNS
B. Azure Application Gateway
C. Azure Private link
D. Azure Traffic Manager

Answer: C

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

NEW QUESTION 46
To complete the sentence, select the appropriate option in the answer area.

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation,

▼
enable the hierarchical namespace.
set Account kind to BlobStorage.
set Performance to Premium.
set Replication to Read-access geo-redundant storage (RA-GRS).

A. Mastered
B. Not Mastered

Answer: A

Explanation:

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation,

▼
enable the hierarchical namespace.
set Account kind to BlobStorage.
set Performance to Premium.
set Replication to Read-access geo-redundant storage (RA-GRS).

NEW QUESTION 50
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
A job that processes sales data once daily is an example of a batch workload.	○	○
A job that calculates a rolling average temperature reading is an example of a streaming workload.	○	○
A job that calculates average revenue per product for the last month is an example of a streaming workload.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Batch processing refers to processing of high volume of data in batch within a specific time span. It processes large volume of data all at once. Batch processing is used when data size is known and finite. It takes little longer time to processes data.
Box 2: Yes
Stream processing refers to processing of continuous stream of data immediately as it is produced. It analyzes streaming data in real time. Stream processing is used when the data size is unknown and infinite and continuous.
Box 3: No

NEW QUESTION 54
To complete the sentence, select the appropriate option in the answer area.

▼	presents content defined by a query.

A heap
A stored procedure
A view
An index

A. Mastered
B. Not Mastered

Answer: A

Explanation:
A View

NEW QUESTION 55
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Stream processing has access to the most recent data received or data within a rolling time window.	○	○
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	○	○
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated

NEW QUESTION 58
To complete the sentence, select the appropriate option in the answer area.

Answer Area

A key/value data store is optimized for

| enforcing constraints |
| simple lookups |
| table joins |
| transactions |

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: simple lookups
A key/value store associates each data value with a unique key. Most key/value stores only support simple query, insert, and delete operations. To modify a value (either partially or completely), an application must overwrite the existing data for the entire value. In most implementations, reading or writing a single value is an atomic operation.
An application can store arbitrary data as a set of values. Any schema information must be provided by the application. The key/value store simply retrieves or stores the value by key.

NEW QUESTION 63
A bank has a system that manages financial transactions.
When transferring money between accounts, the system must never retrieve a value for the source account that reflects the balance before the transfer and a value for the destination account that reflects the balance after the transfer.
Of which ACID semantic is this an example?

A. atomicity
B. isolation
C. durability
D. consistency

Answer: B

Explanation:
Isolation. See course material:

o concurrent transactions cannot interfere with one another, and must result in a consistent database state. o For example, while the transaction to transfer funds from one account to another is in-process, another
transaction that checks the balance of these accounts must return consistent results - the balance-checking transaction can't retrieve a value for one account that reflects the balance before the transfer, and a value for the other account that reflects the balance after the transfer.

NEW QUESTION 68
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
You can copy a dashboard between Microsoft Power BI workspaces.	O	O
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	O	O
A Microsoft Power BY dashboard can display visualizations from a Microsoft Excel workbook.	O	O

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: No
You can duplicate a dashboard. The duplicate ends up in the same Power BI workspace.
There is no current functionality that allows you to move reports from one workspace to another.
Box 2: No
Box 3: Yes

NEW QUESTION 69
Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

A. consistency
B. durability
C. isolation
D. atomicity

Answer: B

NEW QUESTION 72
Your company recently reported sales from the third quarter. You have the chart shown in the following exhibit.

Which type of analysis is shown for the fourth quarter?

A. predictive
B. prescription
C. descriptive
D. diagnostic

Answer: A

NEW QUESTION 76
Which setting can only be configured during the creation of an Azure Cosmos DB account?

A. geo-redundancy
B. multi-region writes
C. production or non-reduction account type
D. API

Answer: D

Explanation:
You must create a separate account for each API.

NEW QUESTION 81
To complete the sentence, select the appropriate option in the answer area.

The massively parallel processing (MPP) engine of Azure
Synapse Analytics ▼

distributes processing across compute nodes.
distributes processing across control nodes.
redirects client connections across compute nodes.
redirects client connections across control nodes.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

The massively parallel processing (MPP) engine of Azure
Synapse Analytics ▼

distributes processing across compute nodes.
distributes processing across control nodes.
redirects client connections across compute nodes.
redirects client connections across control nodes.

NEW QUESTION 86
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory	○	○
Azure Data Factory pipelines can execute other pipelines	○	○
A processing step within an Azure Data Factory pipeline is an activity	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
A pipeline is a logical grouping of activities that together perform a task.
Box 2: Yes
You can construct pipeline hierarchies with data factory.
Box 3: Yes
A pipeline is a logical grouping of activities that together perform a task.

NEW QUESTION 90
To complete the sentence, select the appropriate option in the answer area.

The Azure Cosmos DB [▼] API enables the use of SELECT statements to

| Core (SQL) |
| Gremlin |
| MongoDB |
| Table |

retrieve documents from Azure Cosmos DB.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Azure Cosmos DB SQL API accounts provide support for querying items using the Structured Query Language (SQL) syntax.

NEW QUESTION 92
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
Azure SQL Database includes a fully managed backup service.	○	○
Azure SQL Database has built-in high availability.	○	○
Azure SQL Database can use Azure Advanced Threat Protection (ATP).	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
Azure SQL Database includes a fully managed backup service.	⊡	○
Azure SQL Database has built-in high availability.	⊡	○
Azure SQL Database can use Azure Advanced Threat Protection (ATP).	⊡	○

NEW QUESTION 95
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

A key/value data store is optimized for [simple lookups. ▼]

NEW QUESTION 97
For each of the following statements, select Yes if the statement is true. Otherwise select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory.	○	○
Azure Data Factory pipelines can execute other pipelines.	○	○
A processing step within an Azure Data Factory pipeline is an activity.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
A pipeline is a representation of a data structure within Azure Data Factory.	○	○
Azure Data Factory pipelines can execute other pipelines.	○	○
A processing step within an Azure Data Factory pipeline is an activity.	○	○

NEW QUESTION 101
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
If you have a platform as a service (PaaS) database in Azure, you are responsible for applying operating system updates.	○	○
If you have a platform as a service (PaaS) database in Azure, backups are performed automatically.	○	○
If you have a platform as a service (PaaS) database in Azure, you are responsible for upgrading the database engine.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
Microsoft handles all patching and updating of the SQL and operating system code. You don't have to manage the underlying infrastructure.
Box 2: Yes
SQL Database is a fully managed service that has built-in high availability, backups, and other common maintenance operations.
Box 3: No

NEW QUESTION 102
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
A Microsoft Power BI dashboard is associated with a single workspace.	○	○
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	○	○
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
A Microsoft Power BI dashboard is associated with a single workspace.	○	◉
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	○	◉
A Microsoft Power BI dashboard can display visualizations from a Microsoft Excel workbook.	◉	○

NEW QUESTION 104
Which storage solution supports access control lists (ACLs) at the file and folder level?

A. Azure Data Lake Storage
B. Azure Cosmos DB
C. Azure Queue storage
D. Azure Blob storage

Answer: A

Explanation:
Azure Data Lake Storage Gen2 implements an access control model that supports both Azure role-based access control (Azure RBAC) and POSIX-like access control lists (ACLs).

NEW QUESTION 109
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

In a graph database, edges specify relationships between [tables. ▼]

NEW QUESTION 110
You need to create an Azure Storage account.
Data in the account must replica outside the Azure region automatically.
Which two types of replica can you us for the storage account? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one pint.

A. read-access geo-redundant storage (RA_GRS)
B. zone-redundant storage (ZRS)
C. geo-redundant storage (GRS)
D. locally-redundant storage (LRS)

Answer: AC

Explanation:
‑Azure Storage offers two options for copying your data to a secondary region: Geo-redundant storage (GRS)
‑Geo-zone-redundant storage (GZRS)
With GRS or GZRS, the data in the secondary region isn't available for read or write access unless there is a failover to the secondary region. For read access to the secondary region, configure your storage account to use read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

NEW QUESTION 114

When provisioning an Azure Cosmos DB account, which feature provides redundancy within an Azure region?

A. multi-master replication
B. Availability Zones
C. automatic failover
D. the strong consistency level

Answer: B

Explanation:
With Availability Zone (AZ) support, Azure Cosmos DB will ensure replicas are placed across multiple zones within a given region to provide high availability and resiliency to zonal failures.
Note: Azure Cosmos DB provides high availability in two primary ways. First, Azure Cosmos DB replicates data across regions configured within a Cosmos account. Second, Azure Cosmos DB maintains 4 replicas of data within a region.

NEW QUESTION 115
Select the answer that correctly completes the sentence. See the answer below:

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

An Azure Storage account must have [Hierarchical namespace ▼] enabled to support Azure Data Lake Storage.

NEW QUESTION 118
To complete the sentence, select the appropriate option in the answer area.

A relational database must be used when [▼]
- a dynamic schema is required.
- data will be stored as key/value pairs.
- storing large images and videos.
- strong consistency guarantees are required.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

A relational database must be used when [▼]
- a dynamic schema is required.
- data will be stored as key/value pairs.
- storing large images and videos.
- strong consistency guarantees are required.

NEW QUESTION 123
Select the answer that correctly completes the sentence.

Answer Area

Microsoft SQL Server database tables store	▼

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Microsoft SQL Server database tables store	relational data. ▼

NEW QUESTION 124
You need to recommend a non-relational data store that is optimized for storing and retrieving files, videos, audio stream, and virtual disk images. The data store must store data, some metadata, and a unique ID for each file.
Which type of data store should you recommend?

A. document
B. key/value
C. object
D. columnar

Answer: C

Explanation:
Object storage is optimized for storing and retrieving large binary objects (images, files, video and audio streams, large application data objects and documents, virtual machine disk images). Large data files are also popularly used in this model, for example, delimiter file (CSV), parquet, and ORC. Object stores can manage extremely large amounts of unstructured data.

NEW QUESTION 126
What should you use to automatically delete blobs from Azure Blob Storage?

A. soft delete
B. archive storage
C. the change feed
D. a lifecycle management policy

Answer: D

NEW QUESTION 130
You have an inventory management database that contains the following table.

ProductName	Quantity
Product1	100
Product2	129
Product3	176

Which statement should you use in a SQL query to change the inventory quantity of Product1 to 270?

A. INSERT
B. MERGE
C. UPDATE
D. CREATE

Answer: C

NEW QUESTION 135
Select the answer that correctly completes the sentence.For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
You must apply operating system updates to Azure SQL databases regularly.	○	◉
You need a Microsoft 365 subscription to create an Azure SQL database.	○	○
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
You must apply operating system updates to Azure SQL databases regularly.	○	◉
You need a Microsoft 365 subscription to create an Azure SQL database.	○	◉
You can use existing Microsoft SQL Server licenses to reduce the cost of Azure SQL databases.	◉	○

NEW QUESTION 137
Which statement is an example of Data Manipulation Language (DML)?

A. REVOKE
B. DISABLE
C. CREATE
D. UPDATE

Answer: D

Explanation:
Data Manipulation Language (DML) affect the information stored in the database. Use these statements to insert, update, and change the rows in the database.
BULK INSERT DELETE INSERT SELECT UPDATE MERGE

NEW QUESTION 142
Your company needs to design a database that illustrates the relationships between utilization levels of individual network devices across a local area network.
Which type of data store should you use?

A. graph
B. key/value
C. document
D. columnar

Answer: A

Explanation:
Data as it appears in the real world is naturally connected. Traditional data modeling focuses on defining entities separately and computing their relationships at runtime. While this model has its advantages, highly connected data can be challenging to manage under its constraints.
A graph database approach relies on persisting relationships in the storage layer instead, which leads to highly efficient graph retrieval operations. Azure Cosmos DB's Gremlin API supports the property graph model.

NEW QUESTION 145
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Relational database tables contain columns and rows.	○	○
Indexes in a relational database describe the data types in a table.	○	○
A database view is a virtual table whose content is defined by a query.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Tables are database objects that contain all the data in a database. In tables, data is logically organized in a row-and-column format similar to a spreadsheet. Each row represents a unique record, and each column represents a field in the record.
Box 2: No
An index is an on-disk structure associated with a table or view that speeds retrieval of rows from the table or view.
Box 3: Yes
A view is a virtual table whose contents are defined by a query. Like a table, a view consists of a set of named columns and rows of data.

NEW QUESTION 149
What is a characteristic of non-relational data?

A. a well-defined schema
B. no indexes
C. self-describing entities
D. no unique key values

Answer: D

NEW QUESTION 154
To complete the sentence, select the appropriate option in the answer area.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 156
You have an application that runs on Windows and requires access to a mapped drive. Which Azure service should you use?

A. Azure Files
B. Azure Blob storage
C. Azure Cosmos DB
D. Azure Table storage

Answer: A

Explanation:
Azure Files is Microsoft's easy-to-use cloud file system. Azure file shares can be seamlessly used in Windows and Windows Server.
To use an Azure file share with Windows, you must either mount it, which means assigning it a drive letter or mount point path, or access it via its UNC path.

NEW QUESTION 157
At which level in Azure Cosmos DB can you configure multiple write regions and read regions?

A. database
B. partition
C. collection
D. account

Answer: C

NEW QUESTION 158
You are writing a set of SQL queries that administrators will use to troubleshoot an Azure SQL database. You need to embed documents and query results into a SQL notebook.
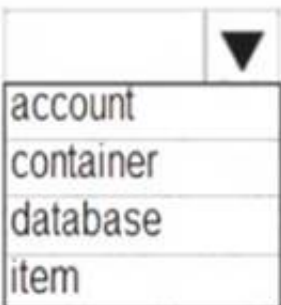What should you use?

A. Microsoft SQL Server Management Studio (SSMS)
B. Azure Data Studio
C. Azure CLI
D. Azure PowerShell

Answer: B

NEW QUESTION 163
To complete the sentence, select the appropriate option in the answer area.

When provisioning an Azure Cosmos DB [▼], you need to specify which type of API you will use.

- account
- container
- database
- item

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Table Description automatically generated

NEW QUESTION 167
Match the types of workloads to the appropriate scenarios.
To answer, drag the appropriate workload type from the column on the left to its scenario on the right. Each workload type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Workload Types		Answer Area	
Batch	Streaming	Workload Type	Analyzing web traffic from the past year
		Workload Type	Classifying images that were uploaded last month
		Workload Type	Tracking how many people are currently using a website

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Batch
Batch processing refers to the processing of blocks of data that have already been stored over a period of time. Box 2: Streaming
Stream processing is a big data technology that allows us to process data in real-time as they arrive and detect conditions within a small period of time from the point of receiving the data. It allows us to feed data into analytics tools as soon as they get generated and get instant analytics results. Box 3: Batch

NEW QUESTION 168
To complete the sentence, select the appropriate option in the answer area.

Answer Area

Relational data uses

| collections |
| columns |
| keys |
| partitions |

to enforce relationships between different tables.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Relational data uses

| collections |
| columns |
| keys |
| partitions |

to enforce relationships between different tables.

NEW QUESTION 173
For each of the following statements, select Yes if the statement is true, Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Azure Synapse Analytics scales storage and compute independently.	○	○
Azure Synapse Analytics can be paused to reduce compute costs.	○	○
An Azure Synapse Analytics data warehouse has a fixed storage capacity.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Azure Synapse Analytics scales storage and compute independently.	○	○
Azure Synapse Analytics can be paused to reduce compute costs.	○	○
An Azure Synapse Analytics data warehouse has a fixed storage capacity.	○	○

NEW QUESTION 177
Select the answer that correctly completes the sentence.

Answer Area

A JSON document is an example of ▼

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

A JSON document is an example of | semi-structured data.   ▼ |

NEW QUESTION 179
You need to design and model a database by using a graphical tool that supports project-oriented offline database development.
What should you use?

A. Microsoft SQL Server Data Tools (SSDT)
B. Microsoft SQL Server Management Studio (SSMS)
C. Azure Databricks
D. Azure Data Studio

Answer: A

NEW QUESTION 184
To complete the sentence, select the appropriate option in the answer area.

| ▼ | physically sorts the data in a table based on the values in a specified column.

| A view |
| A clustered index |
| A stored procedure |
| A nonclustered index |

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated

NEW QUESTION 187
You have a SQL query that combines customer data and order data. The query includes calculated columns. You need to persist the SQL query so that other users can use the query. What should you create?

A. a table
B. an index
C. scalar function
D. a view

Answer: D

Explanation:
A view is a virtual table whose contents are defined by a query. A view acts as a filter on the underlying tables referenced in the view. The query that defines the view can be from one or more tables or from other views in the current or other databases.

NEW QUESTION 191
To complete the sentence, select the appropriate option in the answer area.

By default, each Azure SQL database is protected by | ▼ |

| a network security group (NSG). |
| a server-level firewall. |
| Azure Firewall. |
| Azure Front Door. |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

By default, each Azure SQL database is protected by [▼]

- a network security group (NSG).
- a server-level firewall.
- **Azure Firewall.**
- Azure Front Door.

NEW QUESTION 195

To complete the sentence, select the appropriate option in the answer area.

Answer Area

You can query a graph database in Azure Cosmos DB

- as a JSON document by using a SQL-like language.
- as a partitioned row store by using Cassandra Query Language (CQL).
- as a partitioned row store by using Language-integrated Query (LINQ).
- as nodes and edges by using the Gremlin language.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

You can query a graph database in Azure Cosmos DB

- as a JSON document by using a SQL-like language.
- as a partitioned row store by using Cassandra Query Language (CQL).
- as a partitioned row store by using Language-integrated Query (LINQ).
- as nodes and edges by using the Gremlin language.

NEW QUESTION 197

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Azure Databricks can consume data from Azure SQL Database	○	○
Azure Databricks can consume data from Azure Event Hubs	○	○
Azure Databricks can consume data from Azure Cosmos DB	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Azure Databricks can consume data from SQL Databases using JDBC and from SQL Databases using the Apache Spark connector.
The Apache Spark connector for Azure SQL Database and SQL Server enables these databases to act as input data sources and output data sinks for Apache Spark jobs.
Box 2: Yes
You can stream data into Azure Databricks using Event Hubs. Box 3: Yes
You can run Spark jobs with data stored in Azure Cosmos DB using the Cosmos DB Spark connector. Cosmos can be used for batch and stream processing, and as a serving layer for low latency access.
You can use the connector with Azure Databricks or Azure HDInsight, which provide managed Spark clusters on Azure.

NEW QUESTION 202

Select the answer that correctly completes the sentence.

Answer Area

In a relational database, the [primary key ▼] uniquely identifies each row in a table.

- foreign key
- **primary key**
- WHERE clause
- nonclustered index

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

In a relational database, the [primary key ▼] uniquely identifies each row in a table.

- foreign key
- **primary key**
- WHERE clause
- nonclustered index

NEW QUESTION 203
Match the Azure services to the appropriate locations in the architecture.
To answer, drag the appropriate service from the column on the left to its location on the right. Each service may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Services

Answer Area

- Azure Cognitive Search
- Azure Data Catalog
- Azure Data Factory
- Azure Synapse Analytics

Logs, structured data, and other data structures → **Ingest** Service → **Store** Azure Data Lake Storage → **Preprocess & model** Service → **OLAP** Azure Analysis Services → **Reporting** Microsoft Power BI
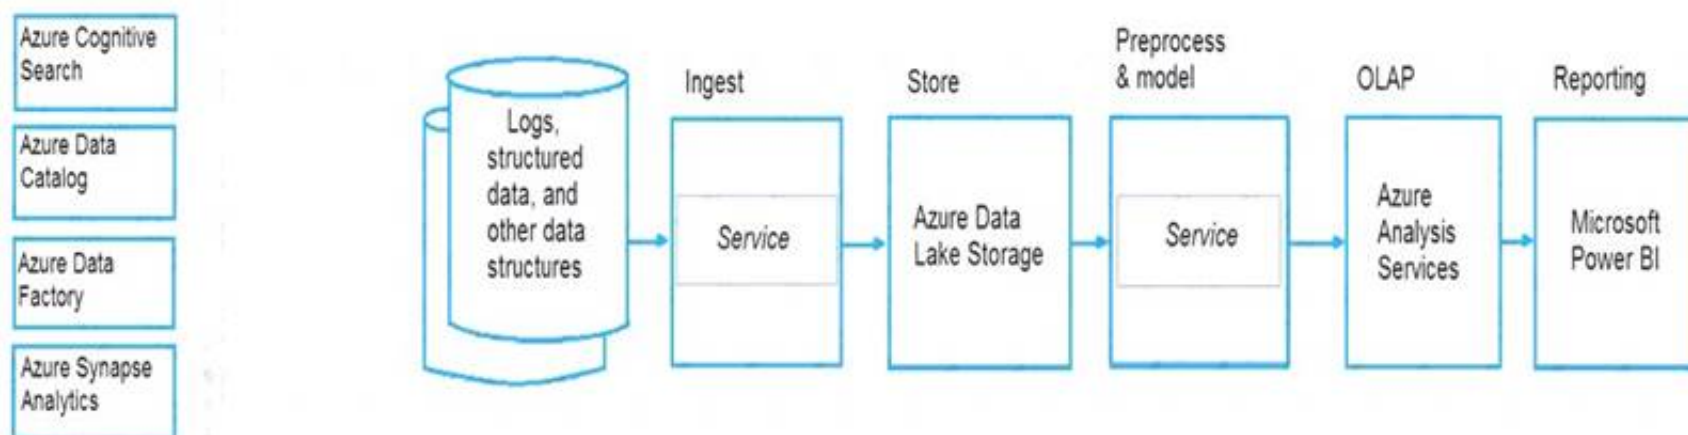
A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box Ingest: Azure Data Factory
You can build a data ingestion pipeline with Azure Data Factory (ADF). Box Preprocess & model: Azure Synapse Analytics
Use Azure Synapse Analytics to preprocess data and deploy machine learning models.

NEW QUESTION 206
Select the answer that correctly completes the sentence.

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a [▼]

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a | graph. | ▼

NEW QUESTION 211
To complete the sentence, select the appropriate option in the answer area.

You can query a graph database in Azure Cosmos DB | ▼ |

| |
as a JSON document by using a SQL-like language.
as a partitioned row store by using Cassandra Query Language (CQL).
as a partitioned row store by using Language-Integrated Query (LINQ).
as nodes and edges by using the Gremlin language.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

You can query a graph database in Azure Cosmos DB | ▼ |

as a JSON document by using a SQL-like language.
as a partitioned row store by using Cassandra Query Language (CQL).
as a partitioned row store by using Language-Integrated Query (LINQ).
as nodes and edges by using the Gremlin language.

NEW QUESTION 216
You need to perform hybrid transactional and analytical processing (HTAP) queries against Azure Cosmos DB data sources by using Azure Synapse Analytics.
What should you use?

A. Synapse pipelines
B. a Synapse SQL pool
C. Synapse Link
D. Synapse Studio

Answer: C

NEW QUESTION 221
What are two benefits of platform as a service (PaaS) relational database offerings in Azure, such as Azure SQL Database? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. complete control over backup and restore processes
B. access to the latest features
C. in-database machine learning services
D. reduced administrative effort for managing the server infrastructure

Answer: BC

Explanation:
Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement.
SQL Database delivers predictable performance with multiple resource types, service tiers, and compute sizes. It provides dynamic scalability with no downtime, built-in intelligent optimization, global scalability and availability, and advanced security options. These capabilities allow you to focus on rapid app development and accelerating your time-to-market, rather than on managing virtual machines and infrastructure.

NEW QUESTION 225
Select the answer that correctly completes the sentence.

Answer Area

The [▼] clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

The [WHERE ▼] clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

NEW QUESTION 228
Select the answer that correctly completes the sentence.

Answer Area

An Azure Storage account must have [▼] enabled to support Azure Data Lake Storage.

Blob soft delete
Hierarchical namespace
Large file shares
Versioning

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Hierarchical namespace
https://learn.microsoft.com/en-us/dynamics365/customer-insights/own-data-lake-storage

NEW QUESTION 229
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	○	○
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	○	○
You can use MySQL Workbench to query Azure Database for MariaDB databases.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	⦿	○
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	⦿	○
You can use MySQL Workbench to query Azure Database for MariaDB databases.	⦿	○

NEW QUESTION 230
What is the primary purpose of a data warehouse?

A. to provide answers to complex queries that rely on data from multiple sources
B. to provide transformation services between source and target data stores
C. to provide read-only storage of relational and non-relational historical data
D. to provide storage for transactional line-of-business (LOB) applications

Answer: C

Explanation:
Consider using a data warehouse when you need to keep historical data separate from the source transaction systems for performance reasons. Data warehouses make it easy to access historical data from multiple locations, by providing a centralized location using common formats, keys, and data models.
Query both relational and nonrelational data.

NEW QUESTION 231
Match the types of workloads the appropriate scenario.
To answer, drag the appropriate workload type from the column on the left to its scenario on the right. Each workload type may be used once, more than once or not at all.
NOTE: Each correct match is worth one point.

Workload Types				Answer Area	
Batch	Micro-batch	Streaming		Workload Type	Data for a product catalog will be loaded every 12 hours to a data warehouse.
				Workload Type	Data for online purchases will be loaded to a data warehouse as the purchases occur.
				Workload Type	Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Workload Types				Answer Area	
Batch	Micro-batch	Streaming		Batch	Data for a product catalog will be loaded every 12 hours to a data warehouse.
				Streaming	Data for online purchases will be loaded to a data warehouse as the purchases occur.
				Micro-batch	Updates to inventory data will be loaded to a data warehouse every 1,000 transactions.

NEW QUESTION 236
Your company has a reporting solution that has paginated reports. The reports query a dimensional model in a data warehouse. Which type of processing does the reporting solution use?

A. stream processing
B. Online Transaction Processing (OLTP)
C. batch processing
D. Online Analytical Processing (OLAP)

Answer: D

NEW QUESTION 241
Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called _____ ▼

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

A block of code that runs in a database is called | a stored procedure. ▼

NEW QUESTION 246
For each of the following statements, select yes if the statement is true, Otherwise, select No. Note: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Batch processing can output data to a file store.	○	○
Batch processing can output data to a relational database	○	○
Batch processing can output data to a NoSQL database.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Batch processing can output data to a file store.	☑	○
Batch processing can output data to a relational database	☑	○
Batch processing can output data to a NoSQL database.	○	☑

NEW QUESTION 248
Your company is designing a data store tor internet-connected temperature sensors. The collected data will be used to analyze temperature trends. Which type of data store should you use?

A. relational
B. columnar
C. graph
D. time series

Answer: D

Explanation:
Time series data is a set of values organized by time. Time series databases typically collect large amounts of data in real time from a large number of sources.

Updates are rare, and deletes are often done as bulk operations. Although the records written to a time-series database are generally small, there are often a large number of records, and total data size can grow rapidly.

NEW QUESTION 250
Match the Azure services to appropriate requirements.
To answer, drag the appropriate services from the column on the left to its requirement on the right. Each service may be used once. more than once, or not at all.
NOTE: Each correct match is worth one point.

A. Mastered
B. Not Mastered

Answer: A

**Explanation:**

NEW QUESTION 255
You need to use Transact-SQL to query files in Azure Data Lake Storage from an Azure Synapse Analytics data warehouse.
What should you use to query the files?

A. Azure Functions
B. Microsoft SQL Server Integration Services (SSIS)
C. PolyBase
D. Azure Data Factory

Answer: C

NEW QUESTION 259
You are reviewing the data model shown in the following exhibit.
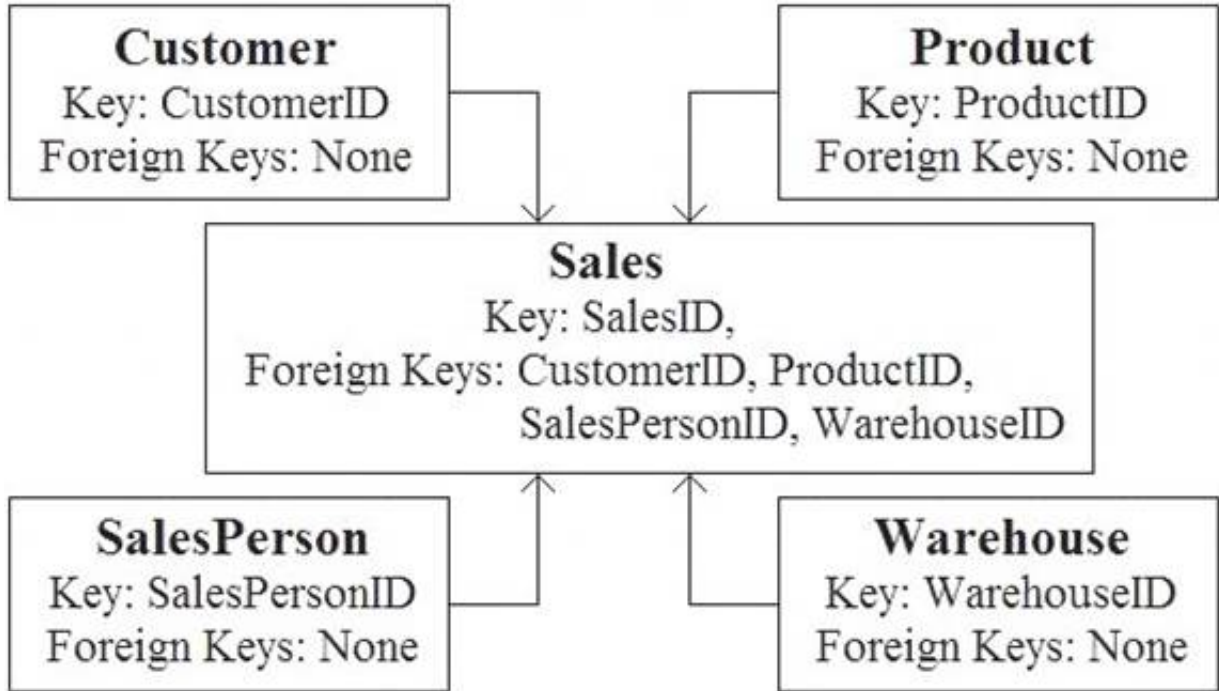
Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point

The data model is a [answer choice].

transactional model
star schema
snowflake schema

Customer is a [answer choice] table.

fact
dimension
bridge

A. Mastered
B. Not Mastered

Answer: A

Explanation:

The data model is a [answer choice].

transactional model
star schema
snowflake schema

Customer is a [answer choice] table.

fact
dimension
bridge

NEW QUESTION 262
Match the terms to the appropriate descriptions.
To answer, drag the appropriate term from the column on the left to its description on the right. Each term may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Terms

Index
View
Table

Answer Area

	A database object that holds data
	A database object whose content is defined by a query
	A database object that helps improve the speed of data retrieval

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Terms

Index
View
Table

Answer Area

Table	A database object that holds data
View	A database object whose content is defined by a query
Index	A database object that helps improve the speed of data retrieval

NEW QUESTION 264

......

Relate Links

100% Pass Your DP-900 Exam with Exambible Prep Materials

https://www.exambible.com/DP-900-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/